

Full-featured peak reduction in right-angled Artin groups

Matthew B. Day

January 2, 2013

Abstract

We prove a new version of the classical peak-reduction theorem for automorphisms of free groups in the setting of right-angled Artin groups. We use this peak-reduction theorem to prove two important corollaries about the action of the automorphism group of a right-angled Artin group A_Γ on the set of k -tuples of conjugacy classes from A_Γ : orbit membership is decidable, and stabilizers are finitely presentable. Further, we explain procedures for checking orbit membership and building presentations of stabilizers. This improves on a previous result of the author's. We overcome a technical difficulty from the previous work by considering infinite generating sets for the automorphism groups. The method also involves a variation on the Hermite normal form for matrices.

1 Introduction

1.1 Background

Let F denote a finite-rank free group with automorphism group $\text{Aut}(F)$. *Peak reduction* is a technique in the study of $\text{Aut}(F)$ that is a key ingredient in the solution of several important problems. J.H.C. Whitehead invented the technique in the 1930's in [15] to provide an algorithm that takes in two conjugacy classes (or more generally, k -tuples of conjugacy classes) from F and determines whether there is an automorphism in $\text{Aut}(F)$ that carries one to the other. In a series of papers in the 1970's [8, 9, 10], McCool used peak-reduction methods to reprove Nielsen's result that $\text{Aut}(F)$ is finitely presented, and to prove that the stabilizer in $\text{Aut}(F_n)$ of a tuple of conjugacy classes in F is also finitely presented.

Given a finite simplicial graph Γ , the *right-angled Artin group* A_Γ is the group given by the finite presentation whose generators are the vertices of Γ , and whose only relations are that two generators commute if and only if they are adjacent as vertices in Γ . These groups are also called “partially commutative groups” and “graph groups”. Free groups are extreme examples of right-angled Artin groups, but this class of groups contains also contains free abelian groups and many other groups. The present paper is concerned with generalizing the peak-reduction method to right-angled Artin groups. In a previous paper [3], the author found a weak generalization of peak reduction and used it to prove that $\text{Aut}(A_\Gamma)$ is finitely presented for every Γ . That paper and its sequel [4] used peak reduction to prove that the stabilizers of certain specific tuples of conjugacy classes in A_Γ are finitely generated, and to investigate an analogue of mapping class groups of surfaces inside $\text{Aut}(A_\Gamma)$. However, the other applications of peak reduction in the free group case do not seem to follow easily from the peak-reduction theorem in [3]. The present paper proves a strong generalization of peak reduction to right-angled Artin groups and uses it to prove two important corollaries of peak reduction.

1.2 Results

First we state our two corollaries of peak reduction. Let A_Γ be a right-angled Artin group.

Theorem 1.1. *There is an algorithm that takes in two tuples U and V of conjugacy classes from A_Γ and either produces an automorphism $\alpha \in \text{Aut}(A_\Gamma)$ with $\alpha \cdot U = V$ or determines that no such automorphism exists.*

Theorem 1.2. *There is an algorithm that takes in a tuple W of conjugacy classes from A_Γ and produces a finite presentation for its stabilizer $\text{Aut}(A_\Gamma)_W$.*

These results follow from our result on peak reduction, which will take a few definitions to state. Let X denote the vertex set of Γ . The length $|u|$ of an element $u \in A_\Gamma$ is the usual length in terms of the generating set X . The length $|w|$ of a conjugacy class w is the minimum length of a representative, and the length $|W|$ of a tuple W of conjugacy classes is the sum of the lengths of its entries.

For $a \in X$, the *star* $\text{st}(a)$ is the subset of X consisting of a and all vertices adjacent to a . For a and b in X , a and b are in the same *adjacent domination equivalence class* if $\text{st}(a) = \text{st}(b)$. We denote the adjacent domination equivalence class of a by $[a]$; if $b \in [a]$ with $a \neq b$, then a is necessarily

adjacent to b . It is entirely possible that $[a] = \{a\}$. We define the set of *generalized Whitehead automorphisms with multiplier set $[a]$* , denoted $\Omega_{[a]}$, to be the subgroup of automorphisms in $\text{Aut}(A_\Gamma)$ such that $\alpha \in \Omega_{[a]}$ if and only if

- for each $b \in X \setminus [a]$, there are $u, v \in \langle [a] \rangle$ with $\alpha(b) = ubv$ and
- for each $b \in [a]$, we have $\alpha(b) \in \langle [a] \rangle$.

In general each $\Omega_{[a]}$ is an infinite group. We define the *permutation automorphisms P* to be the finite subgroup of $\text{Aut}(A_\Gamma)$ with α in P if and only if α restricts to a permutation of $X \cup X^{-1}$. We define the set of *generalized Whitehead automorphisms Ω* to be the union of P with all the $\Omega_{[a]}$ as a varies over the vertices of Γ . It follows from a result of Laurence [6] (see Theorem 2.1 below) that Ω is a generating set for $\text{Aut}(A_\Gamma)$.

Theorem 1.3 (Peak reduction). *Suppose $\alpha \in \text{Aut}(A_\Gamma)$ and W is a tuple of conjugacy classes in A_Γ . Then there is a factorization*

$$\alpha = \beta_m \beta_{m-1} \cdots \beta_2 \beta_1$$

with β_1, \dots, β_m in Ω , such that the sequence of intermediate lengths

$$k \mapsto |\beta_k \cdots \beta_1 \cdot W| \quad \text{for } k = 0, \dots, m$$

strictly decreases from $k = 0$ to $k = k_1$ for some k_1 with $0 \leq k_1 \leq m$, stays constant from $k = k_1$ to $k = k_2$ for some k_2 with $k_1 \leq k_2 \leq m$, and strictly increases from $k = k_2$ to $k = m$. Further, there is an algorithm to find such a factorization.

A major difference between this formulation and the one for free groups is that the set Ω is finite in that setting, but is usually infinite in this setting. Since the classic applications rely on this set being finite, we need some additional results to get the applications to work.

Proposition 1.4. *There is an algorithm that takes in two tuples U and V of conjugacy classes from A_Γ , and a vertex a in Γ , and produces an automorphism α in $\Omega_{[a]}$ with $\alpha \cdot U = V$ or determines that no such automorphism exists. In particular, it is possible to determine whether a tuple U can be shortened by an automorphism from $\Omega_{[a]}$.*

Proposition 1.5. *There is an algorithm that takes in a tuple U of conjugacy classes from A_Γ and a vertex a in Γ , and produces a finite presentation for the stabilizer $(\Omega_{[a]})_U$.*

It turns out that the groups $\Omega_{[a]}$ embed in integer general linear groups, and the key to these propositions is a modified version of the Hermite normal form for integer matrices.

The following example motivates the use of the infinite generating sets.

Example 1.6. This is Example 4.1 from Day [3], which shows that no finite generating set will work for a peak-reduction theorem formulated for the entirety of $\text{Aut}(A_\Gamma)$. In this example, Γ is the four-vertex path graph with vertices a, b, c and d in that order. In [3, Proposition C], we show that for $k \in \mathbb{Z}$, the conjugacy class $w = ad^k$ in this A_Γ is fixed by an automorphism α with $\alpha(a) = ac^k$ and $\alpha(d) = c^{-1}d$ and with α fixing b and c . Further, we show that there is no peak-reduced factorization of α with respect to w by automorphisms that are simpler in a specific sense. This contradicts the existence of a certain formulation of peak-reduction theorem, because such a theorem could be used to produce such factorizations with automorphisms taken from a fixed finite set; however, such a set cannot work for all choices of k . But since the automorphism α is in our set $\Omega_{[c]}$ for any k , this example does not contradict Theorem 1.3.

1.3 Organization of the paper

We postpone the proofs of the more technical results to later sections in the paper. Section 2 contains preliminary facts about right-angled Artin groups. We reduce Propositions 1.4 and 1.5 to propositions about linear groups in Section 3, and we prove our two main application theorems in Section 4. Section 5 proves the propositions about linear groups and we finally prove the peak reduction theorem in Section 6.

2 Preliminaries

2.1 Combinatorial group theory of right-angled Artin groups

The survey by Charney [1] is a good general reference for right-angled Artin groups. We recall a few facts here that are important for understanding the paper. As mentioned above, the *length* of an element u of A_Γ is always the minimal length of an expression for u as a product of generators from Γ and their inverses. Servatius [13] points out that a word in the generators representing u is of minimal length if and only if it is *graphically reduced*, meaning that there is no subword xvx^{-1} where x is a generator or inverse generator and v is a word in those generator that commute with x . Further, Servatius shows that it is possible to get between any two minimal-length

representatives for a word by a sequence of commutation moves. Normal forms found by VanWyk [14] and by Hermiller–Meier [5] are a convenient way to check whether two words represent the same element. Algorithms for checking conjugacy have been described by Liu, Wrathall and Zeger [7], and by Wrathall [16].

One important result is the centralizer theorem of Servatius, from [13]. It states that for any $u \in A_\Gamma$, we can find an element v with $u' = vuv^{-1}$ cyclically reduced (minimal length in its conjugacy class), and for u' cyclically reduced, its centralizer is exactly what one would guess. Specifically, u' can be written as a graphically reduced product $t_1^{k_1} \cdots t_m^{k_m}$ where each t_i is not a proper power and every vertex of Γ used in writing t_i commutes with every vertex of Γ used in writing t_j , for each $i \neq j$. Then the centralizer of u' is generated by the $\{t_i\}_i$ together with each vertex in Γ that commutes with all the vertices appearing in u' . The centralizer of the original u is of course the conjugate of the centralizer of u' by the same conjugating element v .

2.2 Generating Automorphisms

If a and b are vertices in Γ and every vertex that is adjacent to b commutes with a , we say that a *dominates* b . This can happen whether or not a is adjacent to b . Domination is important because it determines whether certain maps defined on generators actually extend to automorphisms. The following definitions are originally from Servatius [13]:

- If a and b are distinct vertices in Γ and a dominates b , then there is an automorphism in $\text{Aut}(A_\Gamma)$ sending b to ba and fixing all other vertices in Γ . This is a *dominated transvection*. If a is not adjacent to b , then there is a different automorphism sending b to ab and fixing all other generators; this is also a dominated transvection. When we need to distinguish between these, we refer to *right* and *left* dominated transvections.
- If a is a vertex in Γ and Y is a connected component of $\Gamma \setminus \text{st}(a)$, then there is an automorphism in $\text{Aut}(A_\Gamma)$ sending each c in Y to aca^{-1} and fixing all other vertices in Γ . This is a *partial conjugation*.
- If π is a symmetry of the graph Γ , then there is an automorphism in $\text{Aut}(A_\Gamma)$ sending each generator x to its image $\pi(x)$. This is called a *graphic automorphism*.
- For each vertex a in Γ , there is an automorphism sending a to a^{-1} and fixing all other generators. This is the *inversion* in a .

If a is adjacent to b then the right and left dominated transvections are the same. Dominated transvections generalize Nielsen moves and also those elementary matrices with a single nonzero off-diagonal entry of one. It is well known that transvections and inversions generate automorphism groups of free groups and general linear groups over the integers, but for general right-angled Artin groups, there are nontrivial partial conjugations and graphic automorphisms that cannot be expressed as products of inversions and transvections. However, these four kinds of automorphisms generate:

Theorem 2.1 (Laurence [6]). *For any Γ , the group $\text{Aut}(A_\Gamma)$ is generated by the finite set consisting of all dominated transvections, partial conjugations, graphic automorphisms, and inversions.*

2.3 Labeled directed graphs

A *labeled directed graph* is a directed multigraph, with self-loops allowed, so that each directed edge carries a label from a pre-specified label set. In this paper, the label sets are always subsets of groups, and we use the convention that a directed edge from a vertex v_1 to a vertex v_2 with label g is also a directed edge from v_2 to v_1 with label g^{-1} . Since all directed edges are considered directed in both directions (with different labels) this means that edge paths are the same as in the underlying undirected graph. In particular, the connected components of a labeled directed graph are the same as the connected components of the underlying undirected graph.

Again we emphasize that in this paper, edge labels are always elements of a pre-specified group. By the *composition of edge labels* of an edge path p in a labeled directed graph, we mean the following: for each edge e on p , we record the label of e in S if the orientations of e and p agree, and if the orientations disagree, we record the inverse of the label of e ; the composition of edge labels is then the composition of these labels and inverses that we recorded, in the order that p traverses them. Or since we consider reversed edges to be labeled with the inverse group element, we can take the composition of edge labels of p simply to be the composition of the labels on the edges in p , interpreted with orientation agreeing with p .

One note: we compose automorphisms like functions, but we compose edge labels on paths in the opposite order, like the usual convention for fundamental groups. This means that if Δ is a directed graph with labels in an automorphism group, then the composition of edge labels along a path labeled with $\alpha_1\alpha_2\cdots\alpha_k$ in Δ is $\alpha_k\cdots\alpha_2\alpha_1$.

Let S be a set of elements of an arbitrary group G and let H be a subgroup of G . The *Schreier graph* of H in G with respect to S is the

labeled, directed multigraph whose vertex set is the set of left cosets of H in G and with an edge labeled by c in S from coset aH to bH if and only if $caH = bH$. Note that since we are not supposing that $S \cup H$ generates G , it will not generally be the case that these Schreier graphs are connected.

Lemma 2.2. *Suppose S is a set of elements of G and Δ is the Schreier graph of H in G with respect to S . Suppose p is an edge path in Δ , starting at a vertex aH and ending at a vertex bH . Let c be the composition of edge labels along p . Then $caH = bH$.*

The proof is an induction argument using the definitions, and is omitted.

In our arguments we need to be able to promote a presentation from a finite-index subgroup of a group to the entire group. Although it is common knowledge that this is possible, we could not find a reference on how to do it. So for completeness, we provide an argument.

Lemma 2.3. *Suppose H is a finite-index subgroup of a group G and we are given a finite presentation $H = \langle S_H | R_H \rangle$ and a finite set $S_G \subset G$ such that the Schreier graph Δ of H in G with respect to $S_H \cup S_G$ is connected. Suppose we are also given an explicit description of Δ . Then we can write down a finite presentation for G .*

Proof. Let p_1, \dots, p_k be generators for the fundamental group of Δ based at H . Since Δ is a finite graph, this is a finite set (generators can be found by selecting a maximal subtree of Δ). Let u_i be the composition of edge labels along p_i for each i ; each u_i is a word in $S_H \cup S_G$. By Lemma 2.2, we know that $u_i \in H$. Let w_i be a word in S_H representing the same element as u_i . Let R_G be the set of relations $\{u_1 w_1^{-1}, \dots, u_k w_k^{-1}\}$. Let $G' = \langle S_H \cup S_G | R_H \cup R_G \rangle$.

Let $\phi: G' \rightarrow G$ send each generator to the element of the same name. Since ϕ sends relators to the identity, it is a well defined group homomorphism. We claim that it is an isomorphism. The group G' acts on Δ through ϕ ; since $S_G \subset G$ is in the image of ϕ , this action is transitive. Let K be the stabilizer of the vertex H of Δ in G' . We note that Δ is then isomorphic to the Schreier graph of K in G' , and therefore the fundamental group of Δ based at H surjects to K by reading off edge labels. In particular, K is generated by the elements u_1, \dots, u_k . Since G' contains relations R_G turning the generators u_1, \dots, u_k into words in S_H , it follows that K is the subgroup of G' generated by S_H . The map ϕ restricts to a map $K \rightarrow G$; since K fixes the vertex H in Δ , the image of this map is in H . We have a map $H \rightarrow G'$ sending each generator the generator of the same name;

since H fixes the vertex H in Δ , the image of this map lies in K . So we have natural maps $H \rightarrow K$ and $K \rightarrow H$; since these are defined by sending generators to generators of the same name, they are inverses of each other. This implies that ϕ is injective: anything in the kernel acts trivially on H and is therefore in K , but $G' \rightarrow G$ restricts to an isomorphism on K .

Now we claim that ϕ is surjective. If not, there is some $g \in G$ not in $\phi(G')$. But G' acts transitively on Δ , so there is $g' \in G'$ with $g' \cdot H = gH$. Then $\phi(g')g^{-1}$ is not in $\phi(G')$, however, it fixes H and is therefore in H . But ϕ maps K maps surjectively to H , a contradiction. \square

3 Reduction to linear problems

3.1 Structure of generalized Whitehead automorphisms

In this section we fix a vertex a and consider the set $\Omega_{[a]}$ of generalized Whitehead automorphisms with multiplying set $[a]$. As defined in the introduction, this is the subgroup of $\text{Aut}(A_\Gamma)$ consisting of automorphisms that multiply elements not in $[a]$ on the right and left by elements of $[a]$, and send elements of $[a]$ to products of elements of $[a]$. Let $\text{dom}(a)$ denote the set of vertices that a dominates, including a itself.

We define $Z_{[a]}$ to be the free abelian group generated by the following symbols:

- for each vertex b in $\text{st}(a) \cap \text{dom}(a)$, a generator r_b ,
- for each vertex b in $\text{dom}(a) \setminus \text{st}(a)$, two generators r_b and l_b , and
- for each connected component Y of $\Gamma \setminus \text{st}(a)$ with at least two vertices, a generator r_Y .

We define a map $\eta: \Omega_{[a]} \rightarrow \text{Aut}(Z_{[a]})$ as follows, given $\alpha \in \Omega_{[a]}$:

- For b in $\text{st}(a) \cap \text{dom}(a)$ and c in $[a]$, the coefficient of r_c in $\eta(\alpha)(r_b)$ is the sum exponent of c in $\alpha(b)$, if $b \notin [a]$ then the coefficient of r_b in $\eta(\alpha)(r_b)$ is one, and the coefficients of all other generators in $\eta(\alpha)(r_b)$ are zero.
- For b in $\text{dom}(a) \setminus \text{st}(a)$ and c in $[a]$, the coefficient of r_c in $\eta(\alpha)(r_b)$ is the sum exponent of c in v and the coefficient of r_c in $\eta(\alpha)(l_b)$ is the sum exponent of c in u , where $\alpha(b) = ubv$; the coefficient of r_b in $\eta(\alpha)(r_b)$ is one, the coefficient of l_b in $\eta(\alpha)(l_b)$ is one, and coefficients of all other basis elements in $\eta(\alpha)(r_b)$ and $\eta(\alpha)(l_b)$ are zero.

- For Y a connected component of $\Gamma \setminus \text{st}(a)$ with at least two vertices and c in $[a]$, the coefficient of r_c in $\eta(\alpha)(r_Y)$ is the sum exponent of c in u where $\alpha(x) = u^{-1}xu$ for some x in Y ; the coefficient of r_Y in $\eta(\alpha)(r_Y)$ is one and the coefficients of all other basis elements in $\eta(\alpha)(r_Y)$ are zero.

We use $\text{GL}(n, R)$ to denote the general linear group of invertible $n \times n$ matrices with entries in a ring R , and we use $M_{n,k}(A)$ to denote the abelian group of $n \times k$ matrices with entries in an abelian group A .

Lemma 3.1. *The map η is a well defined injective homomorphism.*

Further, we can describe its image. Let $n = |[a]|$ and let $k = \dim(Z_{[a]}) - n$. Under an appropriate identification of $Z_{[a]}$ with \mathbb{Z}^{n+k} , the image of η is the set of block-upper-triangular matrices of the form

$$\begin{pmatrix} A & B \\ O & I \end{pmatrix},$$

where A is in $\text{GL}(n, \mathbb{Z})$, B is in $M_{n,k}(\mathbb{Z})$, O is the zero matrix and I is the $k \times k$ identity matrix. Here the matrix A records the coefficients of r_b for $b \in [a]$ and the matrix B records the remaining coefficients.

Proof. First we clarify the statement of the lemma. Let $[a] = \{a_1, \dots, a_n\}$. To identify $Z_{[a]}$ with \mathbb{Z}^{n+k} , we use an identification in which the basis elements r_{a_1}, \dots, r_{a_n} are identified with the first n basis elements of \mathbb{Z}^{n+k} , and the remaining basis elements of $Z_{[a]}$ are identified with the remaining basis elements of \mathbb{Z}^{n+k} . This allows us to identify $\eta(\Omega_{[a]}) \subset \text{Aut}(Z_{[a]})$, assuming that it is well defined.

Since $\alpha \in \Omega_{[a]}$ sends each element b of $[a]$ to some $\alpha(b) \in \langle [a] \rangle$, and the elements of $[a]$ commute, the image of $\eta(\alpha)(r_b)$ is well defined for each $b \in [a]$. For each $b \in \text{dom}(a) \cap \text{st}(a) \setminus [a]$, we know that $\alpha(b)$ is bu for some $u \in \langle [a] \rangle$, since we can commute elements of $[a]$ to the right side of b . Since elements of $[a]$ commute, again $\eta(\alpha)(r_b)$ is well defined. For each $b \in \text{dom}(a) \setminus \text{st}(a)$, we know $\alpha(b) = ubv$ for some $u, v \in \langle [a] \rangle$. These elements u and v are well defined because we cannot commute elements of $[a]$ across this b . Again the coefficients are well defined because $\langle [a] \rangle$ is abelian.

Now consider Y a connected component of $\Gamma \setminus \text{st}(a)$ with at least two vertices. For any b in Y there is some c in Y such that b commutes with c . We know $\alpha(b) = u_1bv_1$ and $\alpha(c) = u_2cv_2$, with $u_i, v_i \in \langle [a] \rangle$. It must be the case that $\alpha(b)$ commutes with $\alpha(c)$. We consider the centralizer of $\alpha(b)$, as indicated by the Servatius centralizer theorem (see Section 2.1). If

$u_1 \neq v_1^{-1}$, then the cyclically reduced form of $\alpha(b)$ contains b and elements of $[a]$. This means that there is an element of $\langle [a] \rangle$ that conjugates $\alpha(c)$ into $\langle \text{st}(b) \cap \text{st}(a) \rangle$. However, this is impossible—any such conjugate of $\alpha(c)$ would have c in it, which is not in $\text{st}(a)$. So this implies that $u_1 = v_1^{-1}$, and of course a parallel argument implies that $u_2 = v_2^{-1}$. Further, the centralizer theorem then implies that $u_1 = u_2$. Repeating this argument on all adjacent pairs in Y , we see that there is a fixed $u \in \langle [a] \rangle$, depending only on α , such that $\alpha(b) = u^{-1}bu$ for any b in Y . Then since $\langle [a] \rangle$ is abelian, the image $\eta(\alpha)(r_Y)$ is well defined.

It is straightforward computation to check that η is a homomorphism.

To see that η is injective and has the specified image, we construct an inverse map θ . Really there is only one way to do this—given a matrix M in the specified image, we read off the first n entries in each column and multiply together elements of $[a]$ with these exponents to get $n + k$ elements of $\langle [a] \rangle$. For each basis element, let u_b be the element determined by the column for r_b , let v_b be determined by the column for l_b , and let u_Y be determined by the column for r_Y . We construct an automorphism $\theta(M)$ that sends $b \in [b]$ to u_b ; $\theta(M)$ sends $b \in \text{dom}(a) \cap \text{st}(a) \setminus [a]$ to bu_b ; $\theta(M)$ sends $b \in \text{dom}(a) \setminus \text{st}(a)$ to v_bbu_b ; and $\theta(M)$ sends each c in a connected component Y with at least two vertices to $u_Y^{-1}cu_Y$. In fact this completely specifies $\theta(M)$ on generators; $\theta(M)$ extends to an endomorphism of A_Γ because the images of generators still satisfy the defining relations for A_Γ , and $\theta(M)$ specifies an automorphism of A_Γ because $\theta(M^{-1})$ is its inverse. Since these maps are clearly inverses of each other, this finishes the proof. \square

In light of Lemma 3.1, we define the group G_1 to be $\text{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\mathbb{Z})$ and we state versions of Propositions 1.4 and 1.5 in this setting. The reason for the subscript 1 will be clear later; we will consider other versions of this group where the subscript denotes a common denominator for certain rational matrix entries.

Proposition 3.2. *There is an algorithm that takes in matrices A and B in $M_{n,k}(\mathbb{Z})$, and returns a matrix D in G_1 with $DA = B$, or determines that there is no such matrix.*

Proposition 3.3. *Suppose A is in $M_{n,k}(\mathbb{Z})$. Then there is an algorithm that produces a finite presentation for the stabilizer of A in G_1 .*

The proofs of these propositions are postponed to Section 5.

3.2 Algorithms for groups of generalized Whitehead automorphisms

In this section we prove Propositions 1.4 and 1.5 modulo Propositions 3.2 and 3.3, which we prove later. We again fix $a \in \Gamma$. The following definition is used repeatedly in this section and also later in the paper.

Definition 3.4. A *syllable* in A_Γ with respect to a is a graphically reduced product of the form $cud \in A_\Gamma$, where u is an element of $\langle \text{st}(a) \rangle$, and c and d are elements of $(\Gamma \setminus \text{st}(a))^{\pm 1}$ or $c = d = 1$. It is a *linear syllable* if $c \neq 1$ and $d \neq 1$, and a *cyclic syllable* if $c = d = 1$. If cud is a linear syllable, then c and d are the *endpoints*.

If w is a conjugacy class in A_Γ , a *decomposition of w into syllables with respect to a* is either of the following:

- if the conjugacy class w has a representative element u in $\langle \text{st}(a) \rangle$, then u itself is a cyclic syllable, and the singleton (u) is a decomposition of w into syllables;
- if w does not have a representative in $\langle \text{st}(a) \rangle$, then a decomposition of w into syllables is a k -tuple of linear syllables for some k :

$$(c_1 u_1 c_2, c_2 u_2 c_3, \dots, c_k u_k c_1),$$

such that the product

$$c_1 u_1 c_2 u_2 c_3 u_3 \cdots c_k u_k \in A_\Gamma$$

is a graphically reduced and cyclically reduced product representing the class w .

The products u or $c_1 u_1 \cdots c_k u_k$ above are the *representatives associated with the decomposition*.

If $W = (w_1, \dots, w_k)$ is a k -tuple of conjugacy classes in A_Γ , a *decomposition of W into syllables with respect to a* is the concatenation, into a single tuple, of some choice of decompositions of the w_i into syllables. Given a decomposition of W into syllables, the *representative associated to the decomposition* is the tuple of elements of A_Γ consisting of the representatives associated with the decompositions of the w_i .

If w is a conjugacy class in A_Γ , there are two possible kinds ambiguities that come up in decomposing w into syllables. First of all, if T is a decomposition of w into syllables, then any cyclic permutation of T is also a

decomposition of w into syllables. Second, if we take the representative associated with a decomposition and get a different graphically reduced product by commuting some letters across syllable boundaries, the resulting product will be associated with a different decomposition of w . These two kinds of ambiguities also come up in decomposing tuples of conjugacy classes into syllables. As we will see below, this second kind of ambiguity is not very important.

The purpose of decomposing things into syllables with respect to a is that it gives us a convenient way to encode the action of $\Omega_{[a]}$ and to keep track of its effect on lengths. We recall the free abelian group $Z_{[a]}$ from Section 3.1. We define a map ν to $Z_{[a]}$ from the set of syllables in A_Γ with respect to a :

- If u is a cyclic syllable then $\nu(u)$ is $\sum_{b \in \text{st}(a) \cap \text{dom}(a)} n_b r_b$, where n_b is the sum exponent of b in u for each $b \in \text{st}(a) \cap \text{dom}(a)$.
- If cud is a linear syllable then $\nu(cud)$ is $v_c + v_d + \sum_{b \in \text{st}(a) \cap \text{dom}(a)} n_b r_b$, where n_b is the sum exponent of b in u for each $b \in \text{st}(a) \cap \text{dom}(a)$, and
 - if c is a positive generator and $c \in \text{dom}(a)$, then $v_c = r_c$,
 - if c^{-1} is a positive generator and $c^{-1} \in \text{dom}(a)$, then $v_c = -l_c$,
 - otherwise c or c^{-1} is in Y , a connected component of $\Gamma \setminus \text{st}(a)$ with two or more vertices, and $v_c = r_Y$;
 - if d is a positive generator and $d \in \text{dom}(a)$, then $v_d = l_c$,
 - if d^{-1} is a positive generator and $d^{-1} \in \text{dom}(a)$, then $v_d = -r_d$, and
 - otherwise d or d^{-1} is in Y , a connected component of $\Gamma \setminus \text{st}(a)$ with two or more vertices, and $v_d = -r_Y$.

The map ν extends diagonally to define a map ν from tuples of syllables to tuples of vectors of $Z_{[a]}$. Note that ν makes no record of the letters in cud from $\text{st}(a) \setminus \text{dom}(a)$. In fact, the map ν gets rid of one of the ambiguities possible in decomposing a class into syllables.

Lemma 3.5. *Suppose w is a conjugacy class in A_Γ and T and T' are decompositions of w into syllables with respect to a . Then the tuple $\nu(T)$ is a cyclic permutation of the entries of the tuple $\nu(T')$. Further, if W is a tuple of conjugacy classes and T and T' are decompositions of W , then $\nu(T)$ is a permutation of $\nu(T')$.*

Proof. As mentioned in Section 2.1, any graphically reduced word representing a given element can be transformed into any other by a sequence of commutations. Further, any graphically and cyclically reduced word representing a conjugacy class can be transformed into any other by a sequence of commutations and cyclic permutations. This observation has a corollary for decompositions of w into syllables. Any decomposition $T = (c_1 u_1 c_2, \dots, c_k u_k c_1)$ of w into syllables can be turned into any other by a sequence of the following moves:

- If some $u_q = u'_q u''_q$ (possibly with $u'_q = 1$ or $u''_q = 1$) and for some $p \leq q$ we have that c_p commutes with c_q , u'_q , u_p , and all c_i and u_i for $i = p+1, \dots, q-1$, and none of the intervening c_i are equal to $c_p^{\pm 1}$, then we can replace the syllable $c_{p-1} u_{p-1} c_p$ with $c_{p-1} u_{p-1} u_p c_{p+1}$, delete $c_p u_p c_{p+1}$ from the list, and break $c_q u_q c_{q+1}$ into two adjacent syllables $c_q u'_q c_p$ and $c_p u''_q c_{q+1}$. In the case that $p = q$, we simply replace $c_{p-1} u_{p-1} c_p$ with $c_{p-1} u_{p-1} u'_p c_p$ and replace $c_p u_p c_{p+1}$ with $c_p u''_p c_{p+1}$.
- The previous move can be done in reverse.
- If some $u_p = u'_p x u''_p$ and for some $q > p$ we have $u_q = u'_q u''_q$ (possibly with any of u'_p, u''_p, u'_q, u''_q equal to 1) such that x commutes with u'_p, c_q, u'_q and all c_i and u_i for $i = p+1, \dots, q-1$, then we can replace the syllable $c_p u_p c_{p+1}$ with $c_p u'_p u''_p c_{p+1}$ and replace $c_q u_q c_{q+1}$ with $c_q u'_q x u''_q c_{q+1}$.
- The previous move can be done in reverse.
- We can cyclically permute the entries of T .

One issue that needs to be addressed with these moves is the fact that they send a well formed syllable decomposition to a well formed syllable decomposition. This follows from the fact that syllables are graphically reduced products. If one of the replacements above results in a syllable that is not graphically reduced, for example, where the c_i cancels with the c_{i+1} , then the original syllable was not graphically reduced. By virtue of moving things only across things they commute with, if a cancellation is possible after the replacement, it was also possible in the first place.

We consider the effect the first move has on $\nu(T)$ in the case that $p < q$. This replaces $c_{p-1} u_{p-1} c_p$ with $c_{p-1} u_{p-1} u_p c_{p+1}$. Since c_p commutes with u_p , but c_p is not adjacent to a , no generator from $\text{st}(a) \cap \text{dom}(a)$ appears in u_p . Since c_p commutes with c_{p+1} , and neither is adjacent to a , this means a dominates neither one and they are in the same component Y of

$\Gamma \setminus \text{st}(a)$. These facts imply that $\nu(c_{p-1}u_{p-1}c_p) = \nu(c_{p-1}u_{p-1}u_p c_{p+1})$. The syllable $c_p u_p c_{p+1}$ is deleted, but since c_p and c_{p+1} are both in Y and since no generator from $\text{st}(a) \cap \text{dom}(a)$ appears in u_p , we have $\nu(c_p u_p c_{p+1}) = 0$. The intervening syllables $c_i u_i c_{i+1}$ for $i = p+1, \dots, q-1$ commute entirely with c_p ; this implies for each i that u_i contains no generators from $\text{st}(a) \cap \text{dom}(a)$ and that both c_i and c_{i+1} are in Y ; this implies that for each i , $\nu(c_i u_i c_{i+1}) = 0$. So although a syllable is deleted before this sequence, thus shifting this sequence forward, there is no effect because this is a sequence of zeros. The syllable $c_q u_q c_{q+1}$ is broken into $c_q u'_q c_p$ and $c_p u''_q c_{q+1}$. For the same reason as the intervening syllables, we see $\nu(c_q u'_q c_p) = 0$. Similarly, c_q and c_p must both be in Y , and u'_q makes no contribution, so $\nu(c_q u_q c_{q+1}) = \nu(c_p u''_q c_{q+1})$. So the effect of the move on $\nu(T)$ is to replace the entry in position p with the same entry, replace zero entries from position $p+1$ to position $q-1$ with zero entries, and replace the entry in position q with an identical entry. Therefore the first move does not affect $\nu(T)$ if $p < q$.

If $p = q$, the first move replaces $c_{p-1}u_{p-1}c_p$ with $c_{p-1}u_{p-1}u'_p c_p$ and replaces $c_p u_p c_{p+1}$ with $c_p u''_p c_{p+1}$. Since u'_p commutes with c_p it contains no generators from $\text{st}(a) \cap \text{dom}(a)$, and we deduce that $\nu(c_{p-1}u_{p-1}c_p) = \nu(c_{p-1}u_{p-1}u'_p c_p)$ and $\nu(c_p u_p c_{p+1}) = \nu(c_p u''_p c_{p+1})$. So the move has no effect on $\nu(T)$ in this case as well.

Next we consider the second kind of move. We replace $c_p u_p c_{p+1} = c_p u'_p x u''_p c_{p+1}$ with $c_p u'_p u''_p c_{p+1}$; since x commutes with c_{p+1} , it is not in $\text{st}(a) \cap \text{dom}(a)$, and therefore $\nu(c_p u_p c_{p+1}) = \nu(c_p u'_p u''_p c_{p+1})$. We also replace $c_q u_q c_{q+1} = c_q u'_q u''_q c_{q+1}$ with $c_q u'_q x u''_q c_{q+1}$; for the same reason, $\nu(c_q u_q c_{q+1}) = \nu(c_q u'_q x u''_q c_{q+1})$. Therefore $\nu(T)$ is unchanged for this kind of move. Of course, if a kind of move leaves $\nu(T)$ unchanged, then the same kind of move in reverse will also leave it unchanged. So the moves above can only change $\nu(T)$ by cyclically permuting its entries.

If W is a tuple of cyclic words and T is a decomposition of W into syllables, then of course, we can cyclically permute the decompositions of the entries in W . Since T is a concatenation of these decompositions of entries, it is possible that different decompositions of W will differ by more complicated permutations—ones that cyclically permute the segments coming from different entries of W . \square

Definition 3.6. Fix a vertex a in Γ . We define an action of $\Omega_{[a]}$ on the set of syllables with respect to a as follows: if $\alpha \in \Omega_{[a]}$ and $s = cud$ where $u \in \langle \text{st}(a) \rangle$ and $c, d \in (\Gamma \setminus [a])^{\pm 1}$ and $\alpha(c) = w_1 c w_2$ and $\alpha(d) = w_3 d w_4$, then define

$$\alpha \cdot s = c w_2 \alpha(u) w_3 d.$$

We extend this action diagonally to the set of k -tuples of syllables for each k .

This action loses the element that α places on the left of the left endpoint of the syllable, and loses the element that α places on the right of the right endpoint of the syllable. We recall the map $\eta: \Omega_{[a]} \rightarrow \text{Aut}(Z_{[a]})$ from Section 3.1. Note that the action of $\Omega_{[a]}$ on syllables just described is the same as the action of $\eta(\Omega_{[a]})$ on $Z_{[a]}$, in that the injective map η intertwines the two actions.

Let $\{a_1, \dots, a_n\} = [a]$. We note that any syllable with respect to a may be written without loss of generality in the form $ca_1^{p_1} \cdots a_n^{p_n}ud$, where $u \in \langle \text{st}(a) \setminus [a] \rangle$. This is because in a general syllable $cu'd$, everything from $[a]$ in u' can be commuted to the left, since $u \in \langle \text{st}(a) \rangle$. The following result is a kind of equivariance between the actions of $\Omega_{[a]}$ on syllables and on $Z_{[a]}$.

Lemma 3.7. *Suppose $ca_1^{p_1} \cdots a_n^{p_n}ud$ is a syllable of A_Γ with respect to a , and suppose $\alpha \in \Omega_{[a]}$. Let p'_i be the coefficient of r_{a_i} in $\alpha \cdot \nu(ca_1^{p_1} \cdots a_n^{p_n}ud)$. Then there are elements $v_1, v_2 \in \langle [a] \rangle$ with*

$$\alpha(ca_1^{p_1} \cdots a_n^{p_n}ud) = v_1ca_1^{p'_1} \cdots a_n^{p'_n}udv_2.$$

Proof. This is a computational exercise in the definitions. \square

Lemma 3.8. *Suppose W is a tuple of cyclic words and*

$$T = (c_1a_1^{p_{1,1}} \cdots a_n^{p_{1,n}}u_1d_1, \dots, c_ma_1^{p_{m,1}} \cdots a_n^{p_{m,n}}u_md_m)$$

is a decomposition of W into syllables, with each $u_i \in \langle \text{st}(a) \setminus [a] \rangle$. Suppose $\alpha \in \Omega_{[a]}$. Let $p'_{i,j}$ be the coefficient on r_{a_j} in the i th coordinate of $\eta(\alpha) \cdot \nu(T)$, for $i = 1, \dots, m$ and $j = 1, \dots, n$. Then

$$\alpha \cdot T = (c_1a_1^{p'_{1,1}} \cdots a_n^{p'_{1,n}}u_1d_1, \dots, c_ma_1^{p'_{m,1}} \cdots a_n^{p'_{m,n}}u_md_m),$$

and $\alpha \cdot T$ is a syllable decomposition of $\alpha \cdot W$.

Proof. This follows from repeated application of Lemma 3.7. In the case that w is a cyclic word and $(c_1u_1c_2, \dots, c_mu_m c_1)$ is its decomposition, we rewrite the representative associated to this decomposition as

$$c_1u_1c_2 \cdot c_2^{-1} \cdot c_2u_2c_3 \cdot c_3^{-1} \cdots c_mu_m c_1 \cdot c_1^{-1}.$$

We apply α to each factor in this product separately. Lemma 3.7 applies to each syllable factor in the product. When we act on a standalone c_i^{-1}

factor in the product, we cancel away the leading and trailing elements from $\langle [a] \rangle$ from the preceding and following syllables (the v_1 and v_2 elements from Lemma 3.7). From the resulting product, in which the exponents on the a_i -factors inside the syllables are from the action of $\eta(\alpha)$ (as in Lemma 3.7), it is easy to see that the corresponding syllable decomposition T' from the lemma statement is a decomposition of $\alpha \cdot w$. The case that W is a tuple of cyclic words follows by applying the same argument to each entry in the tuple separately. \square

Example 3.9. This example illustrates why we need to take a little care with the algorithms for Proposition 1.4 and Proposition 1.5. Suppose for this example that Γ is the graph with four vertices $\{a, b, c, d\}$, with an edge from a to b and an edge from c to d (so A_Γ is $\mathbb{Z}^2 * \mathbb{Z}^2$). Consider the conjugacy classes u and v represented by $cacbc$ and $cbcabcb$ respectively. Choosing syllable decompositions with respect to a arbitrarily, we might choose $T = (cac, cbc, cbc)$ for u and $T' = (cbc, cab, cbc)$ for v . The group $Z_{[a]}$ is generated by r_a, r_b and r_Y , where $Y = \{c, d\}$, and $\nu(T) = (r_a, r_b, r_b)$ and $\nu(T') = (r_b, r_b + r_a, r_b)$. To check whether $\nu(T)$ and $\nu(T')$ are in the same orbit, we apply the algorithm from Proposition 3.2, after choosing an appropriate identification between $Z_{[a]}$ and \mathbb{Z}^3 . We find that $\nu(T)$ and $\nu(T')$ are not in the same orbit. However, $T'' = (cab, cbc, cbc)$ is also a syllable decomposition for v with respect to a , and it is not hard to see that $\nu(T)$ and $\nu(T'')$ are in the same orbit under $\eta(\Omega_{[a]})$: the automorphism sending a to ab^{-1} and fixing the other generators sends one to the other. This automorphism also sends u to v . This example illustrates the need to consider permutations of a syllable decomposition, instead of only considering a single arbitrary decomposition.

Now consider the conjugacy class u represented by $cacb$ in the same group. One syllable decomposition for the conjugacy class is $T = (cac, cbc)$. The automorphism α sending a to b and b to a and fixing the other generators is in $\Omega_{[a]}$, and $\alpha \cdot u = u$. However, $\eta(\alpha) \cdot \nu(T) \neq \nu(T)$. This illustrates the possibility of automorphisms fixing a conjugacy class but not a particular syllable decomposition of that class.

For our finite presentation result, we need refined versions of Propositions 1.4 and 1.5. Specifically, we need to perform the algorithms in these propositions while respecting certain restrictions on the support of automorphisms, which we now define.

Definition 3.10. The *support* of a generalized Whitehead automorphism $\alpha \in \Omega_{[a]}$ is the subset of $X^{\pm 1}$ with

- for b adjacent to or equal to a , b and b^{-1} are both in $\text{supp}(\alpha)$ if $\alpha(b) \neq b$ and neither b nor b^{-1} is in $\text{supp}(\alpha)$ if $\alpha(b) = b$, and
- for b not adjacent to a with $\alpha(b) = ubv$, $b \in \text{supp}(\alpha)$ if and only if $v \neq 1$ and $b^{-1} \in \text{supp}(\alpha)$ if and only if $u \neq 1$.

For $a \in X$ and $S \subset (X \setminus \text{st}(a))^{\pm 1}$, we define $\Omega_{[a],S}$ to be the subset of $\Omega_{[a]}$ consisting of automorphisms α with $\text{supp}(\alpha) \cap S = \emptyset$.

Lemma 3.11. *Suppose $a \in X$ and $S \subset (X \setminus \text{st}(a))^{\pm 1}$. Then $\Omega_{[a],S}$ is a subgroup of $\Omega_{[a]}$. Its image under $\eta: \Omega_{[a]} \rightarrow \text{GL}(n_k, \mathbb{Z})$ is the subset of matrices of the form given in Lemma 3.1, except that those in the image of $\Omega_{[a],S}$ consist of matrices whose upper-left block B has a column of zeros for each basis element representing a subset intersecting S nontrivially.*

Proof. The assertion that this subset is a subgroup is left as an exercise for the reader. For $\alpha \in \Omega_{[a]}$, the columns of the block matrix B in $\eta(\alpha)$ count the appearances of each element from $[a]$ on the right and left sides of $\alpha(b)$ for $b \notin [a]$, whether a adjacently dominates b , non-adjacently dominates b , or does not dominate b (in which case we consider the entire connected component of $\Gamma \setminus \text{st}(b)$ together). We only ignore b if a is adjacent to b and does not dominate it, in which case it follows from Laurence's theorem that α fixes b . The lemma follows from this observation about the columns of B . \square

Lemma 3.12. *Using Proposition 3.2, we can check whether two matrices in $M_{n+k+l,m}(\mathbb{Z})$ are in the same orbit under the group G of block matrices of the form*

$$\begin{pmatrix} A & B & O_{n,l} \\ O_{k,n} & I_k & O_{k,l} \\ O_{l,n} & O_{l,k} & I_l \end{pmatrix},$$

where $A \in \text{GL}(n, \mathbb{Z})$, $B \in M_{n,k}(\mathbb{Z})$, the matrices $O_{i,j}$ are zero matrices and the matrices I_i are identity matrices. Further, one can find a presentation for the stabilizer of a matrix in $M_{n+k+l,m}$ under the action of G using Proposition 3.3.

Proof. Suppose C and D are two matrices in $M_{n+k+l,m}$. If the last l rows of C do not match the last l rows of D , then they cannot be in the same orbit. So we suppose these last l rows match. Next we consider C' and D' in $M_{n+k,m}$, where each is C or D respectively with the last l rows omitted. The group G above is isomorphic to the group G_1 of Proposition 3.2, by the mapping that omits the last l rows and the last l columns, and C is in $G \cdot D$

if and only if C' is in $G_1 \cdot D'$. Of course, this is exactly what Proposition 3.2 checks.

Similarly, the stabilizer of C in G will be isomorphic to the stabilizer of C' in G_1 , which is what Proposition 3.3 provides us with. \square

So instead of proving Proposition 1.4, we prove the following proposition. Proposition 1.4 is the special case where the set S is empty.

Proposition 3.13. *There is an algorithm that takes in two tuples U and V of conjugacy classes from A_Γ , a vertex a of Γ and a subset S of $(X \setminus \text{st}(a))^{\pm 1}$, and produces an automorphism $\alpha \in \Omega_{[a],S}$ with $\alpha \cdot U = V$, or determines that no such automorphism exists.*

Proof. Suppose U and V are two tuples of conjugacy classes of A_Γ . The first step in the algorithm is to form syllable decompositions T of U and T' of V . If the syllable decompositions T and T' do not have the same number of entries, then U and V are not in the same orbit (this follows from Lemma 3.8, since the decomposition of $\alpha \cdot W$ is the same length as the decomposition of W). So we suppose that T and T' have the same number of entries. We consider all the permutations of the entries of T' and select from these the ones T'_1, \dots, T'_m that are also syllable decompositions of V .

Suppose that

$$T = (c_1 a_1^{p_{1,1}} \cdots a_k^{p_{1,k}} u_1 d_1, \dots, c_m a_1^{p_{m,1}} \cdots a_k^{p_{m,k}} u_m d_m).$$

We fix an r from 1 through m , and suppose

$$T'_r = (c'_1 a_1^{p'_{1,1}} \cdots a_k^{p'_{1,k}} u'_1 d'_1, \dots, c'_m a_1^{p'_{m,1}} \cdots a_k^{p'_{m,k}} u'_m d'_m).$$

We define a tuple \hat{T}_r to be T with the exponents of the a_i replaced by those from T'_r :

$$\hat{T}_r = (c_1 a_1^{p'_{1,1}} \cdots a_k^{p'_{1,k}} u_1 d_1, \dots, c_m a_1^{p'_{m,1}} \cdots a_k^{p'_{m,k}} u_m d_m).$$

At this point, we check whether \hat{T}_r is a decomposition of V (this amounts to finding the representative associated to \hat{T}_r and checking whether it represents the V).

If the answer is yes, we use the algorithm from Proposition 3.2 to check whether there is $A \in \eta(\Omega_{[a]})$ with $A \cdot \nu(T) = \nu(\hat{T}_r)$ and with A in the image of $\Omega_{[a],S}$. By Lemma 3.11 we know that $\eta(\Omega_{[a],S})$ is G_1 with the last l columns (without loss of generality) zeroed out in its upper right block (for some l)

and we use the modification of the algorithm from Proposition 3.2 described in Lemma 3.12 to check for such automorphisms.

If our algorithm finds such a matrix A , let $\alpha \in \Omega_{[a],S}$ map to A . By Lemma 3.8, we have that \hat{T}_r is a representative for $\alpha \cdot U$. Of course, in that case, we have $\alpha \cdot U = V$ and the algorithm returns α .

If we get negative answers (either there is no matrix A as above or \hat{T}_r does not represent V), we increment r and check the next T_r . If we try this for each T_r and none of them have a matrix A as above, we declare that U and V are in different orbits.

To show the correctness of the algorithm, we suppose that we have $\alpha \in \Omega_{[a],S}$ with $\alpha \cdot U = V$. We want to show that the algorithm finds an automorphism sending one tuple to the other. We take the coefficients from $\eta(\alpha) \cdot \nu(T)$ and substitute them into the exponents of T to get a tuple T'' ; by Lemma 3.8, T'' is a syllable decomposition for V . By the argument in Lemma 3.5, T'' will differ from other syllable decompositions for V by a sequence of permutations and commutation of elements not in $\text{dom}(a)$ across each other. This means that one of the T'_r will differ from T'' only by the positions of elements not in $\text{dom}(a)$. In particular, $\hat{T}_r = T''$. This means that the algorithm will catch that \hat{T}_r is a representative for V , and then will catch that there is some $\beta \in \Omega_{[a],S}$ with $\eta(\beta) \cdot \nu(T) = \nu(\hat{T}_r)$ (of course this is true for $\beta = \alpha$, but there is no guarantee that the algorithm will catch the same automorphism). So by the contrapositive, if the algorithm does not catch any such automorphism, then no such automorphism exists. \square

Our next goal is to prove Proposition 1.5. Again, we need a slight refinement of the proposition for our argument for finite presentability of stabilizers. So we prove the following, which implies Proposition 1.5 as a special case.

Proposition 3.14. *There is an algorithm that takes in a tuple U of conjugacy classes from A_Γ , and an element $a \in X$ and a subset $S \subset (X \setminus \text{st}(a))^{\pm 1}$ and returns a presentation for the stabilizer $(\Omega_{[a],S})_U$.*

Proof. Let U be a tuple of conjugacy classes in A_Γ . Let T_1 in $Z_{[a]}^M$ be a syllable decomposition of U with respect to a . Let T_1, \dots, T_m be all the permutations of T_1 that are also syllable decompositions of U . The group $\Omega_{[a],S}$ acts on $Z_{[a]}^M$ and a given $\nu(T_i)$ may or may not be in the orbit of $\nu(T_1)$ under this action. We reorder T_1, \dots, T_m so that the intersection of the $\{\nu(T_i)\}_i$ with the orbit of $\nu(T_1)$ is $\nu(T_1), \dots, \nu(T_k)$ for some k . If $\alpha \in \Omega_{[a],S}$ with $\alpha \cdot U = U$ and $i = 1, \dots, k$, then by Lemma 3.8, the element $\alpha \cdot \nu(T_i)$ is the image under ν of a syllable decomposition of U . Then by Lemma 3.5,

the element $\alpha \cdot \nu(T_1)$ is one of $\nu(T_1), \dots, \nu(T_k)$. Therefore $(\Omega_{[a],S})_U$ acts on the finite set $\{\nu(T_1), \dots, \nu(T_k)\}$, and by construction, it acts transitively. Then $(\Omega_{[a],S})_{\nu(T_1)}$ is a finite index subgroup of $(\Omega_{[a],S})_U$.

The observations we just made make it possible to see the correctness of the following algorithm. First we find a syllable decomposition T_1 of U . Then we enumerate all the permutations of T_1 that are also syllable decompositions of U , say T_1, \dots, T_m . Then we use the algorithm from Proposition 3.2 and its modification in Lemma 3.12 to go through the list $\nu(T_1), \dots, \nu(T_m)$ to determine which of these are in the same $\Omega_{[a],S}$ -orbit as $\nu(T_1)$. By relabeling the list, we assume that these are $\nu(T_1), \dots, \nu(T_k)$. Also, as we use the algorithm to check which $\nu(T_i)$ are in the same orbit, we record an example automorphism in $\Omega_{[a],S}$ that sends $\nu(T_1)$ to $\nu(T_i)$, for each i from 1 to k . Let S_1 denote the set of these automorphisms.

Next we use the algorithm from Proposition 3.3 and its modification from Lemma 3.12 to find a presentation for $(\Omega_{[a],S})_{\nu(T_1)}$. Let S_2 denote the generating set for $(\Omega_{[a],S})_{\nu(T_1)}$ from this presentation. For each $\nu(T_i)$ for i from 1 to k , we check where each element of $S_1 \cup S_2$ sends it (since $(\Omega_{[a],S})_U$ acts on $\{\nu(T_1), \dots, \nu(T_k)\}$, the image will be one of these). We construct the finite graph whose vertices are $\nu(T_1), \dots, \nu(T_k)$ and whose edges are labeled by the elements of $S_1 \cup S_2$, with an edge labeled by $\alpha \in S_1 \cup S_2$ from $\nu(T_i)$ to $\nu(T_j)$ if and only if $\alpha \cdot \nu(T_i) = \nu(T_j)$. It is not hard to see that this graph is isomorphic to the Schreier graph of $(\Omega_{[a],S})_{\nu(T_1)}$ in $(\Omega_{[a],S})_U$ with respect to $S_1 \cup S_2$. Note that by the choice of S_1 , this Schreier graph is connected. Using this Schreier graph together with the presentation for $(\Omega_{[a],S})_{\nu(T_1)}$, we then construct a finite presentation for $(\Omega_{[a],S})_U$ using the procedure in Lemma 2.3. \square

4 Applications

4.1 A useful finite graph

In this section we prove our two applications, Theorem 1.1 and Theorem 1.2, modulo the technical results that we prove in the later sections.

Suppose W is an M -tuple of cyclic words in A_Γ and W is of minimal length in its $\text{Aut}(A_\Gamma)$ -orbit. We construct a directed, labeled multigraph Δ associated to W as follows.

- The vertices of Δ are the set of M -tuples of cyclic words of A_Γ of the same length as W .

- For each pair of vertices W_1 and W_2 , possibly with $W_1 = W_2$, and for each permutation automorphism $\alpha \in P$ with $\alpha \cdot W_1 = W_2$, there is a directed edge from W_1 to W_2 labeled by α .
- For each pair of distinct vertices W_1 and W_2 and each generator $a \in X$, if there is an automorphism in $\Omega_{[a]}$ sending W_1 to W_2 , then there is a directed edge from W_1 to W_2 labeled by some $\alpha \in \Omega_{[a]}$ with $\alpha \cdot W_1 = W_2$ (this involves a choice).
- For each vertex W_1 and each generator $a \in X$, there are edges from W_1 to W_1 labeled by a finite generating set for the stabilizer $(\Omega_{[a]})_{W_1}$ (this also involves choices).

Lemma 4.1. *The graph Δ associated to the minimal tuple W is finite and can be effectively constructed.*

Proof. First of all, Δ has finitely many vertices because there are finitely many tuples of a given length. We construct Δ by finding the required edges and attaching them to the 0-skeleton. There are finitely many permutation automorphisms, so we can explicitly check which ones send which vertices to which vertices. For each generator $a \in X$ and each pair of vertices, we use Proposition 1.4 to check whether there is an automorphism in $\Omega_{[a]}$ sending one to the other, and if there is, we add an edge labeled by such an automorphism (Proposition 1.4 gives us one if one exists). For each generator a and each vertex W_1 , we use Proposition 1.5 to get a finite generating set for the stabilizer $(\Omega_{[a]})_{W_1}$. These last two steps are effective since there are only finitely many generators in X and vertices in Δ . \square

Lemma 4.2. *If α is the composition of edge labels on a path from a vertex W_1 to a vertex W_2 in Δ , then $\alpha \cdot W_1 = W_2$.*

Proof. This is true for paths of length one by construction and true in general by Lemma 2.2. \square

Lemma 4.3. *Suppose W_0 is a vertex in Δ and W_0 is minimal length in its automorphism orbit. If W' is also a vertex in Δ and $\alpha \in \text{Aut}(A_\Gamma)$ with $\alpha \cdot W_0 = W'$, then there is a path p in Δ from W_0 to W' such that the composition of edge labels along p is α .*

Proof. We peak reduce α with respect to W_0 by elements of Ω , which is possible by Theorem 1.3. Suppose $\alpha = \alpha_k \cdots \alpha_1$ is the resulting factorization. Let $W_i = \alpha_i \cdots \alpha_1 \cdot W_0$ for $i = 1, \dots, k$, so that $W_k = W'$. Since W_0 is minimal length, the factorization being peak reduced means that $|W_k| = |W_0|$ for

$i = 0, \dots, k$. Then each W_i is a vertex in Δ . If α_i is a permutation automorphism, then there is an edge from W_{i-1} to W_i labeled by α_i by construction. If $\alpha_i \in \Omega_{[a]}$ for some a , then there is an edge from W_{i-1} to W_i labeled by β_i for some $\beta_i \in \Omega_{[a]}$. Then $\beta_i^{-1}\alpha_i$ stabilizes W_{i-1} . By construction, the edge loops at W_{i-1} contain generators for that stabilizer, and there is a path p_i in the loops at W_{i-1} whose edge labels compose to be $\beta_i^{-1}\alpha_i$. So following p_i and then the edge labeled by β_i gives a path p'_i from W_{i-1} to W_i such that the composition of labels on p'_i is α_i . Composing these paths as i goes from 1 to k gives a path from W_0 to W_k whose edge label composition is α . \square

4.2 Orbit membership and finite generation

Lemma 4.4. *Suppose W is a tuple of conjugacy classes from A_Γ . Then W is minimal length in its $\text{Aut}(A_\Gamma)$ -orbit if and only if it cannot be shortened by any element of $\Omega_{[a]}$ for any $a \in X$.*

Proof. It is clear that a minimal-length tuple cannot be shortened, so we prove the other direction. Suppose for contradiction that there is some W' with $|W'| < |W|$ and some $\alpha \in \text{Aut}(A_\Gamma)$ with $\alpha \cdot W = W'$. We peak reduce α with respect to W by Theorem 1.3. Then $\alpha = \beta_k \cdots \beta_1$ with $\beta_i \in \Omega$ for all i , where $k \mapsto |\beta_k \cdots \beta_1 \cdot W|$ is a sequence of lengths that decreases, stays level, and then increases (with any of these phases possibly omitted). Since $|\alpha \cdot W| < |W|$, the decreasing phase cannot be omitted, and therefore $|\beta_1 \cdot W| < |W|$. This contradiction proves the lemma. \square

Proof of Theorem 1.1. Let U and V be two M -tuples of conjugacy classes from A_Γ ; we want to check whether they are in the same $\text{Aut}(A_\Gamma)$ -orbit. We start by enumerating the tuples of conjugacy classes from A_Γ that are strictly shorter than $|U|$. Of course there are only finitely many such conjugacy classes (even so, this step is a disappointing bottleneck in the algorithm). For each U' strictly shorter than U , and each class $[a]$ with $a \in \Gamma$, we use Proposition 1.4 to check whether U' is in the same orbit as U under $\Omega_{[a]}$. If it is, we replace U by U' and repeat the previous step again, checking whether one of the $\{\Omega_{[a]}\}_a$ can shorten U . We stop when we have verified that none of these sets of automorphisms can shorten U . By Lemma 4.4, this resulting U is of minimal length.

After we shorten U as much as possible, we do the same to V . If the minimal lengths are different, we declare that the orbits are different. Now we suppose that U and V are both minimal length in their automorphism orbits with $|U| = |V|$; let Δ be the graph from Section 4.1 above, constructed

using U (the vertices of Δ are tuples of length $|U|$). At this point, we check whether U and V are in the same connected component of Δ (this is doable since Δ is a finite graph). By Lemma 4.2, if U and V are in the same component, then there is an automorphism $\alpha \in \text{Aut}(A_\Gamma)$ with $\alpha \cdot U = V$, and we can find such an α by composing the edge labels on a path from U to V . Conversely, if there is an automorphism α sending U to V , then there is a path from U to V in Δ and both are in the same connected component. This is true by Lemma 4.3, which uses Theorem 1.3. \square

At this point we can quickly prove an intermediate result. By $\pi_1(\Delta, W)$ we mean the fundamental group of Δ based at W ; this may be interpreted either combinatorially or topologically.

Proposition 4.5. *Suppose W is a tuple of cyclic words in A_Γ . Then there is an algorithm to find a finite generating set for the stabilizer $\text{Aut}(A_\Gamma)_W$.*

Proof. Altering W by an automorphism sends $\text{Aut}(A_\Gamma)_W$ to a conjugate, so we are free to replace W with a representative of its orbit of minimal length. We find such an element using Lemma 4.4 and Proposition 1.4. Now assuming that W has minimal length in its orbit, we construct the graph Δ as above. Then composition of edge labels defines a homomorphism $\pi_1(\Delta, W) \rightarrow \text{Aut}(A_\Gamma)$ (this cannot fail because the domain is free). By Lemma 4.2 (with $W_1 = W_2 = W$), the image of this homomorphism lies in $\text{Aut}(A_\Gamma)_W$. By Lemma 4.3 (with $W_0 = W' = W$), this homomorphism surjects on $\text{Aut}(A_\Gamma)_W$. \square

4.3 Relations

Our next goal is to show that stabilizers of tuples of conjugacy classes are finitely presentable. Before we start, we need to record some relations that hold among the generalized Whitehead automorphisms. First we mention relations between classic Whitehead automorphisms. In our terminology, a *classic Whitehead automorphism* is either

- a permutation automorphism from P (an automorphism that restricts to a permutation of $X^{\pm 1}$) or
- an automorphism α with a special element a of X , called its *multiplier*, such that α sends a to a and sends each $b \in \Gamma \setminus \{a\}$ to one of b , ba , $a^{-1}b$ or $a^{-1}ba$.

The relations between classic Whitehead automorphisms from Definition 2.6 of Day [3] play a limited role in the current paper. We can essentially treat these as a black box.

We also need peak reduction of long-range Whitehead automorphisms, which we likewise can treat as a black box. A generalized Whitehead automorphism α in $\Omega_{[a]}$ is *short-range* if $\alpha(b) = b$ for all b not adjacent to a (but α may do anything to $\text{st}(a)$). It is *long-range* if the restriction of α to $\text{st}(a)^{\pm 1}$ is a permutation of that set (but α may do anything to $X \setminus \text{st}(a)$). More generally, an automorphism in $\text{Aut}(A_\Gamma)$ is *long-range* if and only if it can be factored as a product of long-range dominated transvections, partial conjugations, inversions and graphic automorphisms. The important fact about long-range automorphisms is:

Theorem 4.6 (Day, from [3], Theorem A). *If W is a tuple of conjugacy classes in A_Γ and $\alpha \in \text{Aut}(A_\Gamma)$ is a long-range automorphism, then α has a factorization by classic long-range Whitehead automorphisms and permutation automorphisms that is peak reduced with respect to W (that is, α has a factorization by these kinds of automorphisms that satisfies the conclusions of Theorem 1.3).*

We need the following relations. These are like the Steinberg relations from algebraic K-theory.

Lemma 4.7. *Suppose $a, b \in \Gamma$ with $[a] \neq [b]$, and we have $\alpha \in \Omega_{[a]}$, and $\beta \in \Omega_{[b]}$. Further suppose α restricts to the identity on $[b]$. Then $\alpha\beta\alpha^{-1} \in \Omega_{[b]}$ and $\alpha \cdot \beta \cdot \alpha^{-1} = (\alpha\beta\alpha^{-1})$ is an identity among generalized Whitehead automorphisms if either of the following is true:*

- a is adjacent to b , or
- $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ and β restricts to the identity on $[a]$.

We further note that if β restricts to the identity on $[a]$, then in fact $\alpha\beta\alpha^{-1} = \beta$.

Proof of Lemma 4.7. This follows by straightforward computations, which we describe in broad strokes. Let γ denote $\alpha\beta\alpha^{-1}$ and suppose that c is in X ; we need to show that $\gamma(c)$ is in $\langle [b] \rangle$ if $c \in [b]$, and that $\gamma(c) = u_1cu_2$ with $u_1, u_2 \in \langle [b] \rangle$ if $c \notin [b]$. If $c \in [b]$, then it is clear that $\gamma(c) = \beta(c)$. If $c \in [a]$, then in all three cases it is straightforward to show that $\gamma(c) = ucv$ for some u, v in $\langle [b] \rangle$. The reasons for this are different in the three cases above. Now suppose that $c \notin [a] \cup [b]$. Of course, there are u_1 and u_2 in $\langle [a] \rangle$

with $\alpha(c) = u_1cu_2$ and v_1 and v_2 in $\langle [b] \rangle$ with $\beta(c) = v_1cv_2$. Then $\alpha\beta\alpha^{-1}$ sends c to

$$\alpha\beta\alpha^{-1}(u_1^{-1})v_1u_1cu_2v_2\alpha\beta\alpha^{-1}(u_2^{-1}).$$

In the first case, $\alpha\beta\alpha^{-1}(u_1)$ differs from u_1 by an element of $\langle [b] \rangle$ and u_1 and v_1 commute, so the result follows. In the second case, either u_i or v_i is trivial for $i = 1, 2$ and the result follows. \square

4.4 The stabilizer presentation complex

Let W be an M -tuple of cyclic words in A_Γ that is minimal length in its automorphism orbit. To prove Theorem 1.2, we build a finite cellular 2-complex Z whose fundamental group is the stabilizer $\text{Aut}(A_\Gamma)_W$. The 1-skeleton Z^1 is like the graph Δ defined earlier, but with some extra edges. In order to define a map $\pi_1(Z, W) \rightarrow \text{Aut}(A_\Gamma)$, we give the Z^1 the structure of a labeled multigraph.

- The vertices Z^0 are the M -tuples of cyclic words in A_Γ of the same length as W and in the same orbit.
- For each pair of vertices (not necessarily distinct) W_1 and W_2 and each classic Whitehead automorphism α , we add an edge from W_1 to W_2 labeled by α if $\alpha \cdot W_1 = W_2$. This includes the cases where α is a permutation automorphism.
- For each pair of distinct vertices W_1 and W_2 , each generator $a \in X$ and each subset $S \subset (X \setminus \text{st}(a))^{\pm 1}$, if there is an element of $\Omega_{[a],S}$ sending W_1 to W_2 , then we make sure there is an edge from W_1 to W_2 labeled by some such element.
- For each vertex W_1 , each generator $a \in X$ and each subset $S \subset (X \setminus \text{st}(a))^{\pm 1}$, the labels on edges from W_1 to itself must include a generating set from a presentation for the stabilizer $(\Omega_{[a],S})_{W_1}$.

Like Δ , the graph Z^1 can be effectively constructed using Propositions 3.13 and 3.14. Instead of checking whether a tuple is in the same length as W to decide whether to use it as a vertex, it may be more efficient to construct Δ as above, discard other connected components, and then add extra edges to form Z^1 . Since we only consider vertices in the same orbit as W , Z^1 is automatically connected and each vertex W_1 of Z is minimal length in its orbit.

Next we define the several situations where we add 2-cells to Z . When we say a 2-cell “reads off” a word starting at a given vertex, we mean that

we glue in the 2-cell so that its boundary follows the path whose edge labels form that word.

- (C1) For each vertex W_1 , each generator $a \in X$ and each subset $S \subset (X \setminus \text{st}(a))^{\pm 1}$, the self-edges at W_1 labeled by elements of $\Omega_{[a],S}$ give a generating set for $(\Omega_{[a],S})_{W_1}$ by construction. We add 2-cells reading off the relations between these elements, and we add enough 2-cells so that the subcomplex of Z spanned by these edges forms a presentation complex for $(\Omega_{[a],S})_{W_1}$. This is possible by Proposition 3.14.
- (C2) Suppose there is an edge starting at the vertex W_1 with label α , where α is a long-range generalized Whitehead automorphism but not a classic Whitehead automorphism. We find a path from W_1 to $\alpha \cdot W_1$ with label sequence $\gamma_1 \gamma_2 \cdots \gamma_k$, where each γ_i is a classic long-range Whitehead automorphism, and glue in a 2-cell reading off the difference between these two paths. This is possible by Theorem 4.6: we peak-reduce α with respect to W_1 to get the factorization $\alpha = \gamma_k \cdots \gamma_1$; since the factorization is peak reduced and W_1 is minimal length, each intermediate image $\gamma_i \cdots \gamma_1 \cdot W_1$ is also a vertex of Z and this word defines an edge path.
- (C3) Whenever we find an edge loop in Z whose labels read off one of the relations between classic Whitehead automorphisms from Day [3, Definition 2.6], we add a 2-cell bounding this edge loop. These relations fall into ten classes, are easily recognizable, and each such relation has length at most five.
- (C4) Suppose W_1 is in Z^0 , α is a generalized Whitehead automorphism labeling an edge starting at W_1 , and β is an inner automorphism that is also a classic Whitehead automorphism. Then there is a factorization $\alpha\beta\alpha^{-1} = \gamma_k \cdots \gamma_1$, where the γ_i are also inner classic Whitehead automorphisms. The inner classic Whitehead automorphisms label loops in Z^1 and we glue in a 2-cell reading off the difference between these two factorizations starting at $\alpha \cdot W_1$. We repeat this for each such W_1 , α and β .
- (C5) Suppose W_1 is in Z^0 and there is a closed edge path p starting at W_1 whose edge labels α, β, γ are in $\Omega_{[a]}$ for some $a \in X$. Then $\gamma\beta\alpha \in (\Omega_{[a]})_{W_1}$. Since the labels on the loops at W_1 include generators for $(\Omega_{[a]})_{W_1}$, we know that there is an edge path w consisting of loops at W_1 whose composition represents the same automorphism as $\gamma\beta\alpha$. Then we add a 2-cell to Z whose boundary follows p and then follows

w backwards. We add such a cell for each vertex W_1 on each such path p involving at least two vertices.

- (C6) Suppose W_1 is in Z^0 , $\alpha \in P$ and $\beta \in \Omega_{[b]}$ for some b , and both α and β are edge labels on edges starting at W_1 . Then since α is a permutation automorphism, $\alpha\beta \cdot W_1$ is also a vertex of W_1 . It is easy to see that the element $\alpha\beta\alpha^{-1}$ is in $\Omega_{[\alpha(b)]}$ and sends $\alpha \cdot W_1$ to $\alpha\beta \cdot W_1$. By construction, there is an edge in Z^1 labeled by some $\gamma \in \Omega_{[\alpha(b)]}$ with $\gamma\alpha \cdot W_1 = \alpha\beta \cdot W_1$, and therefore $\alpha\beta^{-1}\alpha^{-1}\gamma$ is in $(\Omega_{[\alpha(b)]})_{W_1}$. Since the loop edge labels at W_1 include a generating set for $(\Omega_{[\alpha(b)]})_{W_1}$, we have a path w in these loops where the composition of edge labels represents $\alpha\beta^{-1}\alpha^{-1}\gamma$. Then we add a 2-cell to Z whose boundary, starting at $\alpha \cdot W_1$, follows α^{-1} then β then α then γ^{-1} , and then w . We repeat this for each vertex W_1 and each such pair α and β .
- (C7) Suppose W_1 is in Z^0 and $a, b \in X$ and $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$ are edge labels on edges starting at W_1 and α and β satisfy the hypotheses of Lemma 4.7 (so $\alpha|_{[b]}$ is the identity, and either a is adjacent to b , or $\beta|_{[a]}$ is also the identity and $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$). Then Proposition 6.7 below implies that $\alpha\beta \cdot W_1$ is the same length as W_1 and therefore is a vertex in Z^0 . So by construction, there is an automorphism $\gamma \in \Omega_{[a]}$ labeling an edge from $\beta \cdot W_1$ to $\alpha\beta \cdot W_1$ (α is such an automorphism, but there is no guarantee that our construction of Z^1 found this particular automorphism.) Further, $\alpha\beta\alpha^{-1}$ is in $\Omega_{[b]}$ by Lemma 4.7, and $\alpha\beta\alpha^{-1}$ sends $\alpha \cdot W_1$ to $\alpha\beta \cdot W_1$. By construction of Z^1 , there is $\delta \in \Omega_{[b]}$ sending $\alpha \cdot W_1$ to $\alpha\beta \cdot W_1$, and there is an edge in Z^1 labeled by δ from $\alpha \cdot W_1$ to $\alpha\beta \cdot W_1$. We note that $\alpha\gamma^{-1}$ fixes $\alpha\beta \cdot W_1$, and therefore there is a path w_1 in the edge loops at $\alpha\beta \cdot W_1$ where the composition of the edge labels represents $\alpha\gamma^{-1}$. Similarly, $\delta\alpha\beta^{-1}\alpha^{-1}$ fixes $\alpha\beta \cdot W_1$ and there is a path w_2 in the edge loops at $\alpha\beta \cdot W_1$ where the composition of these edge loops represents $\delta\alpha\beta^{-1}\alpha^{-1}$. We glue a 2-cell into Z whose boundary, starting at W_1 , follows β , then γ , then w_1 , then w_2 , then δ^{-1} , and finally α^{-1} . We repeat this for each vertex W_1 and each pair α, β at W_1 satisfying the hypotheses of Lemma 4.7.

This completes the construction of Z . We note that in principle, Z can be effectively constructed from Γ , since in each of the cases (C1) through (C7), there are only finitely many cases in which we may have to add a 2-cell.

Lemma 4.8. *Composition of edge labels defines a surjective homomorphism $\pi_1(Z, W) \rightarrow \text{Aut}(A_\Gamma)_W$.*

Proof. The proof of Lemma 4.2 goes through with Δ replaced by Z^1 , since the extra edges in Z^1 still indicate the action by their labels. Then we have a well defined homomorphism $\pi_1(Z^1, W) \rightarrow \text{Aut}(A_\Gamma)_W$ for the same reasons as in the proof of Proposition 4.5. By Lemma 4.3 (with $W_0 = W' = W$), this homomorphism surjects on $\text{Aut}(A_\Gamma)_W$. By the Seifert–Van Kampen theorem, the kernel of the natural map $\pi_1(Z^1, W) \rightarrow \pi_1(Z, W)$ is normally generated by the boundary loops of the 2–cells. By construction, each of these boundary loops maps to the trivial automorphism. Then the homomorphism descends to a homomorphism $\pi_1(Z, W) \rightarrow \text{Aut}(A_\Gamma)_W$, which is necessarily surjective. \square

The following proposition is the key to Theorem 1.2.

Proposition 4.9. *Suppose p is an edge loop in Z based at W that maps to the trivial automorphism. Then p can be homotoped relative to W to an edge loop whose edge labels consist entirely of permutation automorphisms and inner automorphisms.*

Since the proof of Proposition 4.9 uses the structure of the proof of Theorem 1.3, we postpone it to Section 6. This statement is all we need to prove the finite presentation result.

Proof of Theorem 1.2. Suppose W is a tuple of conjugacy classes in A_Γ . First we find a minimal-length representative of the orbit of W using Proposition 1.4 and Lemma 4.4. Of course, replacing W with an image of itself under an automorphism will not change the isomorphism type of $\text{Aut}(A_\Gamma)_W$ —it will only replace it with a corresponding conjugate of itself in $\text{Aut}(A_\Gamma)$. So we replace W with a minimal representative of its orbit.

We construct the complex Z with respect to W as described above, and consider the map $\pi_1(Z, W) \rightarrow \text{Aut}(A_\Gamma)_W$ from Lemma 4.8. If this map is an isomorphism, then the Seifert–Van Kampen theorem implies that the stabilizer is finitely presented, since Z is a finite complex. To prove the theorem, it is enough to show that the map is injective, since it is already surjective by Lemma 4.8.

To show injectivity, we assume we have an edge loop p based at W , such that the composition of the edge labels of p yields the trivial homomorphism. By Proposition 4.9, we assume we have homotoped p to an edge loop whose edge labels are permutation automorphisms and inner automorphisms. We use the 2–cells of type (C6) to slide these inner automorphisms past the permutation automorphisms. Then the inner automorphisms label loops at the base vertex W . The multiplication table of the group P is included in the

relations that the (C3) cells bound, so we can eliminate all the permutation automorphisms from p by homotoping across these cells. Then p reads off a product of inner automorphisms representing the trivial automorphism. We can rewrite any inner automorphism in Ω as a product of inner automorphisms that are also classic Whitehead automorphisms by homotoping across (C1) cells. The group of inner automorphisms of A_Γ is isomorphic to another right-angled Artin group ($A_{\Gamma'}$, where Γ' is Γ with all the vertices representing central generators deleted). Further, this isomorphism carries the inner classic Whitehead automorphisms to the standard generating set of the right-angled Artin group. So any word in the inner classic Whitehead automorphisms that represents the trivial automorphism can be eliminated by applying commutation relations. These commutation relations are given by 2-cells in Z (redundantly as (C3) or (C7) cells), so we can homotope p to the trivial edge path at W . \square

5 Orbits of matrices

In this section we prove Propositions 3.2 and 3.3.

5.1 Partly rational linear problems

We fix $n \geq 1$ and $k \geq 0$ and consider block-upper triangular matrices of the form

$$\begin{pmatrix} A & B \\ O & I \end{pmatrix},$$

with A in $\mathrm{GL}(n, \mathbb{Z})$ and B in $M_{n,k}(\mathbb{Q})$. This is the semidirect product $\mathrm{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\mathbb{Q})$, where $\mathrm{GL}(n, \mathbb{Z})$ acts on $M_{n,k}(\mathbb{Q})$ on the left by multiplication. We will vary the kinds of entries we wish to consider in the upper-right block, so we let $G_{\mathbb{Q}}$ denote $\mathrm{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\mathbb{Q})$, and for a positive integer d , we let G_d denote $\mathrm{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$. We deliberately restrict the coefficients in the upper left-block to \mathbb{Z} , so that G_d will always be a finite-index subgroup in G_1 (we discuss this more below). A key tool in our discussion will be the following modified version of the Hermite normal form for integer row reduction. For more on the Hermite normal form, see Cohen [2], section 2.4.2. Note that Cohen describes Hermite normal form for column reduction, whereas we use Hermite normal form for row reduction.

Definition 5.1. Suppose A is a matrix in $M_{n+k,m}(\mathbb{Q})$. Roughly speaking, A is in $G_{\mathbb{Q}}$ -normal form if

- \mathbb{Q} -linear combinations of rows $n + 1$ through $n + k$ have been added to rows 1 through n to reduce them as much as possible, and
- there is a multiple of the block of rows 1 through n that is a matrix in Hermite normal form for integer row reduction.

Precisely, $A = (a_{ij})$ is in $G_{\mathbb{Q}}$ -normal form if

- for each j , $1 \leq j \leq m$, if there is a linear combination of rows $n + 1$ through $n + k$ that is nonzero in column j but is zero in all previous rows, then every entry in column j from row 1 through row n is zero,
- there is an increasing sequence p_1, \dots, p_l of pivot column positions, for some l with $0 \leq l \leq n$, such that for each i from 1 through l :
 - the entry a_{i,p_i} is positive,
 - the entries $a_{i,1}, \dots, a_{i,p_i-1}$ in row i preceding column p_i are all 0,
 - for each i' , $i' = 1, \dots, i - 1$, the entry a_{i',p_i} satisfies $0 \leq a_{i',p_i} < a_{i,p_i}$ (the entries above the pivot position are nonnegative and less than the pivot), and
- rows from $l + 1$ through n contain only zero entries.

We want to show that every matrix is equivalent to a unique one in this form, and to understand the stabilizer of a matrix in this form. First we prove uniqueness.

Lemma 5.2. *Suppose A and B are matrices in $M_{n+k,m}(\mathbb{Q})$ in $G_{\mathbb{Q}}$ -normal form and Q is in $G_{\mathbb{Q}}$ with $B = QA$. Let d be the smallest positive integer with $Q \in G_d$. Then $Q = Q_1Q_2$ with $Q_2 \in \{I\} \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$ and $Q_1 \in \text{GL}(n, \mathbb{Z}) \ltimes \{O\}$, such that $Q_2A = A$ and such that the first l columns of Q_2 are the same as those of the identity matrix, where l is the number of pivots in A . In particular, $A = B$.*

Proof. First of all, since $G_{\mathbb{Q}}$ cannot alter rows $n + 1$ through $n + k$ by multiplication on the left, we see that the bottom k rows of B and A must be identical. In particular, for each position j , if there is a linear combination of the bottom k rows in A that is trivial in columns 1 through $j - 1$ but nontrivial in column j then there is one for B as well, and vice versa. Then the set of columns that are forced to be trivial by the first condition in the definition of $G_{\mathbb{Q}}$ -normal form is the same in both A and B . The matrix Q can be factored as Q_1Q_2 , where $Q_2 \in \{I\} \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$ and $Q_1 \in \text{GL}(n, \mathbb{Z}) \ltimes \{O\}$.

$\{O\}$, because G_d is an internal semidirect product of these two subgroups. Then Q_2 fixes A , because Q_2 can only change an entry in A that is forced to be zero by the definition, and Q_1 cannot change the set of columns that only have zero entries in their first n rows.

So we have $B = Q_1 A$, with $Q_1 \in \text{GL}(n, \mathbb{Z}) \ltimes \{O\}$. The uniqueness of Hermite normal form implies that $A = B$, since Q_1 only changes the top n rows of A and the top $n \times k$ blocks of A and B are already rational multiples of matrices in Hermite normal form. However, our statement about the form of Q_1 is a little stronger, so we prove the lemma as stated.

Let p_1, \dots, p_l be the pivots of A . We perform induction on the hypothesis that for i with $i \leq l$, columns 1 through $p_{i+1} - 1$ of A and B match and columns 1 through i of Q_1 are the same as those of identity matrix. The hypothesis is true for $i = 0$ since columns 1 through $p_1 - 1$ of A are trivial and the equation $B = Q_1 A$ is only possible if the corresponding columns of B are also trivial. Now we fix an i with $1 \leq i \leq l$ and consider column p_i of A and B . Certainly entries $i + 1$ through n of column p_i of B are zero, since B is in $G_{\mathbb{Q}}$ -normal form and column p_i is left of the $(i + 1)$ st pivot column of B (if it exists). Consider i' with $i < i' \leq n$. By the inductive hypothesis, the first $i - 1$ entries of the i' th row in Q_1 are zero. We dot this row i' with the column p_i in A to get an entry in B that we know to be zero. Since entries $i + 1$ through n of column p_i in A are zero, the only possibly nonzero term in this dot product is the product of the i', i entry of Q_1 with the i, p_i entry of A . So entry i', i of Q_1 is zero, and varying i' , we see that every below-diagonal entry in Q_1 in column i is zero. Next we note that if the diagonal entry i, i of Q_1 were zero, then Q_1 would have determinant zero. So this entry is nonzero, and therefore position i, p_i in B has a nonzero entry and is the pivot there. Since the pivot entries are positive, the fact that Q_1 has determinant ± 1 implies that the i, i entry of Q_1 is 1. Then position i, p_i matches in A and B . If any above-diagonal entry in column i of Q_1 is nonzero, then an entry in column p_i of B above the pivot will not be reduced modulo the pivot entry. This then implies that all the columns before p_{i+1} of A and B match, since these columns have only zero entries in positions $i + 1$ through n .

The induction continues until we reach the last pivot position of A , showing that the first l columns of Q_1 match those of the identity matrix. Since rows $l + 1$ through n of A are zero rows, this is enough to deduce that $B = A$. \square

Now we show existence of matrices in normal form.

Proposition 5.3. *Every matrix A in $M_{n+k,m}(\mathbb{Q})$ is associated to a matrix B in $M_{n+k,m}(\mathbb{Q})$ in $G_{\mathbb{Q}}$ -normal form, with $A = QB$ for some $Q \in G_{\mathbb{Q}}$. The matrix B is unique.*

Proof. We prove existence by supplying an algorithm; of course uniqueness is then the result of Lemma 5.2. The algorithm is a row reduction algorithm. Multiplication on the left by elements of $G_{\mathbb{Q}}$ allows us to replace any of the top n rows of A by itself plus a rational linear combination of the bottom k rows, or to replace any row in the top n by itself plus an integer linear combination of the other top n rows, or to permute the top n rows, or to multiply any of the top n rows by -1 .

The first part of the algorithm is to use the bottom k rows of A to simplify A as much as possible; this is step 1 below. The second part is to perform integer row reduction and reduce the entries above the pivots as much as possible.

Step 1: We start by setting $j = 1$; j is the position of the column we are trying to simplify. We consider the map $\mathbb{Q}^k \rightarrow \mathbb{Q}^{j-1}$ that sends a k -tuple of coefficients to the corresponding linear combination of the bottom k rows of A , restricted to their first $j - 1$ columns. We find generators for the kernel of this map. Each generator gives us a linear combination of the bottom k rows of A that is zero in its first $j - 1$ entries; if some generator's linear combination is nonzero in the j th column, we add rational multiples of this linear combination to the top n rows of A to zero out their j th column entries. Of course this leaves the previous columns unaffected. (In the case that $j = 1$, we simply check whether some row in the bottom k rows has a nonzero entry in its first column, and if so, we use its multiples to zero out the top n entries of the first column.) So we replace A with an equivalent matrix with the first n entries of column j zeroed out if possible. We then replace j with $j + 1$ and repeat this step. We do this until we have tried it for all columns of A .

After completing the previous step, we perform the procedure to turn the top $n \times m$ block of A into a rational multiple of a matrix in Hermite normal form for integer row reduction. Although this is standard, we include it here for completeness.

Step 2: If the top $n \times m$ block of A is now the zero matrix, then A is in $G_{\mathbb{Q}}$ -normal form and we are done. Otherwise we start the second part by setting j to be the first column with a nonzero entry in its first n rows. We initialize our sequence of pivots by setting $l = 0$, so that there are no pivots in the pivots sequence p_1, \dots, p_l .

Step 3: By construction, $j > l$ and all entries in row $l + 1$ through row

n in columns 1 through $j - 1$ are zero. We look the entries of column j from row $l + 1$ through row n , and choose a row i whose column- j entry has minimal nonzero absolute value among these (if all the entries were zero, we would increment j and loop, but the choice of j should prevent this). We add integer multiples of row i to each row from $l + 1$ through n in order to diminish the sum of the absolute values of the entries in column j . We continue this until either another row's j th entry becomes smaller in absolute value than that of row i , or else row i becomes the only entry from position $l + 1$ through n with a nonzero entry. If another row's j th entry becomes smaller in absolute value than that of row i , then we replace i with the position of that row and repeat this step. If position i becomes the only entry in column j from position $l + 1$ through n that is nonzero, then we proceed to the next step.

Step 4: The entry in position i is the unique nonzero entry in column j among rows from $l + 1$ through n . We permute rows $l + 1$ through n of A so that this nonzero entry is now in row $l + 1$. We replace l with $l + 1$ and set the new pivot position p_l to be j . We replace row l with its multiple by ± 1 to ensure that the pivot entry in position l, p_l is positive. We then add integer multiples of row l to rows 1 through $l - 1$ to make sure that these entries are nonnegative and strictly less than the pivot entry. Since all entries to the left of the pivot entry are zero, this does not affect the previous columns of the matrix. We then set j to be the next column with a nonzero entry in rows $l + 1$ through n and return to step 3. If no such column exists, we are done. This finishes the algorithm.

It is clear that this procedure terminates. The loop in step 1 repeats once for each column. After step 1, the least common denominator of the entries of A does not change; call this number d . The loop in step 3 always terminates because we decrease the sum of the absolute values of the entries of column j of A by at least $1/d$ with each iteration. The loop going back from step 4 to step 3 can be repeated at most m times since it requires a new column each time.

It is also clear that the output of this procedure is a matrix in $G_{\mathbb{Q}}$ -normal form. The matrix coming out of step 1 satisfies the first property in the definition: if there were a way to use the bottom k rows to zero out the top n in the j th column without disturbing the previous columns, we would have used it already. This is not disturbed by the remainder of the algorithm, which never changes a column with only zeros in its first n rows to one with a nonzero entry there. The output of the algorithm has its list of pivot columns, and by construction, the pivots satisfy the conditions in the definition. Of course, we only stop producing pivots when all the

remaining rows in the first n are zero rows, so we satisfy the last condition in the definition.

Let A denote the matrix input to the procedure and B the output, which is in $G_{\mathbb{Q}}$ -normal form. Keeping track of the row moves performed in this algorithm and composing them gives us a matrix Q in $G_{\mathbb{Q}}$ with $B = QA$. \square

Now we turn our attention to stabilizers in G_d . We need the following.

Lemma 5.4. *Suppose the group G acts on the group H , G has the presentation $\langle S_G | R_G \rangle$, H has the presentation $\langle S_H | R_H \rangle$, and the set R_C consists of words $gh^{-1}g^{-1}w_{g,h}$ for all $g \in S_G$ and $h \in S_H$, where $w_{g,h}$ is a word in S_H representing ghg^{-1} . Then*

$$\langle S_G \cup S_H | R_G \cup R_H \cup R_C \rangle$$

is a presentation for the semidirect product $G \ltimes H$.

The proof is left as an exercise for the reader.

Proposition 5.5. *Suppose d is a positive integer and A is a matrix in $M_{n+k,m}(\frac{1}{d}\mathbb{Z})$ in $G_{\mathbb{Q}}$ -normal form. Then there is an effective procedure to give a finite presentation for the stabilizer $(G_d)_A$.*

Proof. First we pin down the stabilizer; then we will find a presentation for it. Let l be the number of pivot rows of A . Suppose $Q \in G_d$ and $A = QA$. By Lemma 5.2, we have $Q = Q_1 Q_2$ where $Q_2 \in \{I\} \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$, $Q_1 \in \text{GL}(n, \mathbb{Z}) \ltimes \{O\}$, $Q_2 A = A$ and the first l columns of Q_1 are the same as those of the identity matrix. This tell us that the nontrivial block of Q_1 is in $M_{l,n-l}(\mathbb{Z}) \ltimes \text{GL}(n-l, \mathbb{Z})$, in other words the nontrivial block of Q_1 is itself a block-upper-triangular matrix of the form

$$\begin{pmatrix} I & B \\ O & C \end{pmatrix}$$

where $B \in M_{l,n-l}(\mathbb{Z})$, $C \in \text{GL}(n-l, \mathbb{Z})$, I is the $l \times l$ identity matrix and O is the $(n-l) \times l$ zero matrix. The configuration of blocks implies that $\text{GL}(n-l, \mathbb{Z})$ acts on $M_{l,n-l}(\mathbb{Z})$ on the right, as the semidirect product notation reflects.

The matrix Q_2 is in the stabilizer of A in $\{I\} \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$. Each row of Q_2 acts by adding a rational linear combination of rows $n+1$ through $n+k$ of A to some row of A in the top n , and being in the stabilizer means that each such rational linear combination has trivial value. So each row of the upper-right $n \times k$ block of Q_2 is an element of the kernel of

the group homomorphism $(\frac{1}{d}\mathbb{Z})^k \rightarrow \mathbb{Q}^m$ that sends a k -tuple of coefficients to its corresponding linear combination of rows $n+1$ through $n+k$ of A . Letting $K \subset (\frac{1}{d}\mathbb{Z})^k$ denote this kernel, we see that the stabilizer of A in $\{I\} \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z})$ is isomorphic to K^n .

This makes it easy to see that the stabilizer of A in G_d is

$$(M_{l,n-l}(\mathbb{Z}) \rtimes \mathrm{GL}(n-l, \mathbb{Z})) \ltimes K^n,$$

since each element of this group stabilizes A and any Q stabilizing A is certainly in this group by the above argument.

Since K is the kernel of a map $(\frac{1}{d}\mathbb{Z})^k \rightarrow \mathbb{Q}^m$, it is a finite-rank free abelian group and we can find a basis for K . This in turn yields a basis for K^n , which is also free abelian. Likewise $M_{l,n-l}(\mathbb{Z})$ is a free abelian group with an obvious basis. Of course, this means that K^n and $M_{l,n-l}(\mathbb{Z})$ have obvious finite presentations, where the generators are the given bases and the relations state that all pairs of basis elements commute. The group $\mathrm{GL}(n-l, \mathbb{Z})$ has a generating set given by transvections (elementary matrices with a single nonzero off-diagonal entry of 1) and inversions (matrices sending a single basis element to its inverse and fixing the others). The finite presentation for $\mathrm{SL}(n-l, \mathbb{Z})$ from Milnor [11, Chapter 10], can easily be modified to give a finite presentation for $\mathrm{GL}(n-l, \mathbb{Z})$. The conjugate of a generator of $M_{l,n-l}(\mathbb{Z})$ by a generator of $\mathrm{GL}(n-l, \mathbb{Z})$ can easily be written down as a product of generators of $M_{l,n-l}(\mathbb{Z})$. We can tabulate this data for all choices of pairs of generators. Using Lemma 5.4, this data together with the presentations for $\mathrm{GL}(n-l, \mathbb{Z})$ and $M_{l,n-l}(\mathbb{Z})$ can be combined into a finite presentation for $M_{l,n-l}(\mathbb{Z}) \rtimes \mathrm{GL}(n-l, \mathbb{Z})$. Finally, we can tabulate the action of generators of $M_{l,n-l}(\mathbb{Z}) \rtimes \mathrm{GL}(n-l, \mathbb{Z})$ on generators of K^n (again in terms of generators of K^n). This data, together with the obvious presentation for K^n and the presentation for $M_{l,n-l}(\mathbb{Z}) \rtimes \mathrm{GL}(n-l, \mathbb{Z})$, can be combined to give a presentation for the stabilizer of A in G_d , again using Lemma 5.4. \square

5.2 Integer linear problems

Let d be a fixed positive integer. To exploit our rational results in the previous section, we use a crossed homomorphism to keep track of the cosets of G_d in G_1 . Let $\rho: G_d \rightarrow M_{n,k}(\mathbb{Z}/d\mathbb{Z})$ be the following composition

$$G_d = \mathrm{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\frac{1}{d}\mathbb{Z}) \rightarrow M_{n,k}(\frac{1}{d}\mathbb{Z}) \rightarrow M_{n,k}(\mathbb{Z}) \rightarrow M_{n,k}(\mathbb{Z}/d\mathbb{Z}),$$

where the maps are the second coordinate projection from the semidirect product, then multiplication by d , then reduction modulo d .

Lemma 5.6. *The map ρ is a crossed homomorphism: for $A, B \in G_d$, we have*

$$\rho(AB) = A \cdot \rho(B) + \rho(A),$$

where G_d acts on $M_{n,k}(\mathbb{Z}/d\mathbb{Z})$ via the projection $G_d \rightarrow \mathrm{GL}(n, \mathbb{Z})$ and the standard left action of $\mathrm{GL}(n, \mathbb{Z})$ on $M_{n,k}(\mathbb{Z}/d\mathbb{Z})$.

The kernel of ρ is G_1 and further, the set of preimages of elements of $M_{n,k}(\mathbb{Z}/d\mathbb{Z})$ under ρ is precisely the set of left cosets of G_1 in G_d .

Proof. We consider the definition of the product operation in $\mathrm{GL}(n, \mathbb{Z}) \ltimes M_{n,k}(\mathbb{Z}/d\mathbb{Z})$:

$$(A, B) \cdot (C, D) = (AC, AD + B).$$

Of course this means that the projection $G_d \rightarrow M_{n,k}(\frac{1}{d}\mathbb{Z})$ is a crossed homomorphism with respect to the action via the canonical left action of $\mathrm{GL}(n, \mathbb{Z})$. Since the map $M_{n,k}(\frac{1}{d}\mathbb{Z}) \rightarrow M_{n,k}(\mathbb{Z}/d\mathbb{Z})$ is equivariant with respect to this action, the map ρ is a crossed homomorphism.

If $A \in G_d$ is in the kernel of ρ , this means that the entries in its upper-right block are divisible by d after being multiplied by d , in other words that they are integers. So the kernel of ρ is G_1 . Now suppose $B \in M_{n,k}(\mathbb{Z}/d\mathbb{Z})$. We can pick representatives for the residue classes of the entries of B to get an element \tilde{B} in $M_{n,k}(\mathbb{Z})$ mapping to B . Then ρ maps $(I, \frac{1}{d}\tilde{B})$ in G_d to B . We claim that $\rho^{-1}(B)$ is the coset $(I, \frac{1}{d}\tilde{B}) \cdot G_1$. If (C, D) is in G_1 , then

$$\rho((I, \frac{1}{d}\tilde{B})(C, D)) = \rho(\frac{1}{d}\tilde{B}) + \rho(D) = B + 0.$$

This implies that the coset is a subset of the preimage. On the other hand, if $\rho((C, D)) = B$, then $D - \frac{1}{d}\tilde{B}$ is in $M_{n,k}(\mathbb{Z})$, and $(C, D - \frac{1}{d}\tilde{B}) \in G_1$ with $(C, D) = (I, \frac{1}{d}\tilde{B})(C, D - \frac{1}{d}\tilde{B})$. In other words, the preimage is a subset of the coset. This shows that every preimage is a coset.

Now suppose that $(A, B)G_1$ is a coset. Let $\overline{B} = \rho((A, B))$. If $(C, D) \in G_1$, then $\rho((A, B)(C, D)) = B \cdot \rho(D) + \rho(B) = 0 + \overline{B}$, so this coset is a subset of this preimage. If $(C, D) \in G_d$ and $\rho((C, D)) = \overline{B}$, then $(C, D) = (A, B)(A^{-1}C, A^{-1}(D - B))$ with $(A^{-1}C, A^{-1}(D - B)) \in G_1$, so this preimage is a subset of this coset. So every coset is a preimage. \square

Now we can prove our proposition on determining orbit membership under the action of G_1 on $M_{n+k,m}(\mathbb{Z})$.

Proof of Proposition 3.2. Let A and B be in $M_{n+k,m}(\mathbb{Z})$; we wish to find a matrix $D \in G_1$ with $DA = B$, or show that no such matrix exists. We start by computing the $G_{\mathbb{Q}}$ -normal forms of A and B , using the algorithm

in Proposition 5.3. If these normal forms are different, then A and B are in distinct $G_{\mathbb{Q}}$ -orbits (by uniqueness of the normal form, Lemma 5.2). Certainly if A and B are in different $G_{\mathbb{Q}}$ -orbits, then they are also in different G_1 -orbits. So we suppose that A and B have the same $G_{\mathbb{Q}}$ -normal form $N \in M_{n+k,m}(\mathbb{Q})$. Suppose $Q, R \in G_{\mathbb{Q}}$ with $N = QA$ and $N = RB$. Let d be the least common denominator of the entries of N , Q and R . Let S be the generating set from the presentation for the stabilizer of N in G_d from Proposition 5.5. Let Δ be the Schreier graph of G_1 in G_d with respect to S .

Claim. *A and B are in the same G_1 -orbit if and only if the vertices QG_1 and RG_1 are in the same connected component of Δ .*

First we suppose that QG_1 and RG_1 are in the same connected component of Δ . Let $C \in G_d$ be the composition of edge labels on an edge path from QG_1 to RG_1 . Then $CQG_1 = RG_1$ by Lemma 2.2. In particular, $R^{-1}CQ$ is in G_1 , and $R^{-1}CQ$ sends A to B .

Conversely, we suppose that there is D in G_1 with $DA = B$. Then RDQ^{-1} fixes N . By Proposition 5.5, RDQ^{-1} is a word in S . Starting at QG_1 , we form an edge path in Δ by following the edges labeled by this expression for RDQ^{-1} . This is possible and unambiguous since Δ , being a Schreier graph, has exactly one edge with each label entering and leaving each vertex. Then by Lemma 2.2, the terminus of the path is $RDQ^{-1}QG_1 = RG_1$. This proves the claim.

The algorithm should then be clear at this point. First we compute the $G_{\mathbb{Q}}$ -normal forms of A and B and report that A and B are in different G_1 -orbits if these $G_{\mathbb{Q}}$ -normal forms differ. If the normal forms are the same matrix N , we find matrices Q and R with $QA = N$ and $RB = N$ and find the lowest common denominator d of the entries of Q , R and N . Then we find a generating set S for the stabilizer of N in G_d and construct the Schreier graph Δ of G_1 in G_d with respect to S . This is possible since Δ is finite by Lemma 5.6. The crossed homomorphism ρ gives a convenient way to construct Δ : $CDG_1 = EG_1$ if and only if $\rho(CD) = \rho(E)$, if and only if $C\rho(D) + \rho(C) = \rho(E)$ for any $D, E \in G_d$ and $C \in S$. Next in the algorithm, we check whether QG_1 and RG_1 are in the same connected component. If not, we report that A and B are in different G_1 -orbits. If they are in the same connected component, we take C to be the composition of edge labels along a path from QG_1 to RG_1 , and report that $R^{-1}CQ$ is a matrix in G_1 sending A to B . \square

Now we find a presentation for the stabilizer in G_1 of a matrix A in $M_{n+k,m}(\mathbb{Z})$.

Proof of Proposition 3.3. Let N be the $G_{\mathbb{Q}}$ -normal form of A and let $Q \in G_{\mathbb{Q}}$ be an element with $N = QA$, as found using Proposition 5.3. Let d be the least common denominator of the entries of N and Q . Let S be a generating set for the stabilizer of N in G_d , as given by Proposition 5.5. Let Δ be the Schreier graph of G_1 in G_d with respect to S . Let S' be $\{Q^{-1}CQ | C \in S\}$; note that $S' \subset (G_d)_A$. Finally, let Δ' be the Schreier graph of $(G_1)_A$ in $(G_d)_A$ with respect to S' . The proof of the proposition will follow from the claim:

Claim. *As a directed multigraph, Δ' is isomorphic to the connected component of $Q^{-1}G_1$ in Δ , by an isomorphism that sends edge labels in S' to their conjugates by Q in S .*

To prove the claim, we start by defining a map on vertices from Δ' to Δ . Let $B \in (G_d)_A$. We send the vertex $B \cdot (G_1)_A$ of Δ' to the vertex $Q^{-1}BG_1$ of Δ . This map is well defined: if C is also in $(G_d)_A$ with $B \cdot (G_1)_A = C \cdot (G_1)_A$, then $B^{-1}C \in (G_1)_A \subset G_1$ and therefore $Q^{-1}BG_1 = Q^{-1}CG_1$. Conversely, the map is injective: given $B, C \in (G_d)_A$ with $Q^{-1}BG_1 = Q^{-1}CG_1$, we see that $B^{-1}C \in G_1$; since $(G_1)_A = (G_d)_A \cap G_1$, this means that $B^{-1}C \in (G_1)_A$ and therefore that $B \cdot (G_1)_A = C \cdot (G_1)_A$. Suppose there is an edge labeled by C in S' from $B \cdot (G_1)_A$ to $B' \cdot (G_1)_A$ in Δ' . Then there is an edge labeled by $Q^{-1}CQ$ in S from $Q^{-1}BG_1$ to $Q^{-1}B'G_1$. This is immediate from the definition of the Schreier graph. The reverse implication also holds, so the map is an isomorphism of directed multigraphs and respects labels as described.

All that is left in the claim is to show that Δ' is connected. Suppose $B \cdot (G_1)_A$ is a vertex of Δ' . Of course this means that $B \in (G_d)_A$. Then $QBQ^{-1} \in (G_d)_N$, and by the definition of S , QBQ^{-1} can be expressed as a product of elements of S . Of course, this means that B can be expressed as a product of elements of S' . Since Δ' is a Schreier graph, we can trace out a unique edge path starting at $(G_1)_A$ using the labels from the given expression for B as a product of elements of S' . By Lemma 2.2, the terminus of this path is $B \cdot (G_1)_A$. Since B was arbitrary, this means that Δ' is connected. This proves the claim.

Now we use the claim to finish the proposition. Let Z be the presentation 2-complex for $(G_d)_N$ using the finite presentation from Proposition 5.5. Of course, the generating set for this presentation is S . The inner automorphism of G_d given by conjugating by Q^{-1} sends $(G_d)_N$ to $(G_d)_A$ and sends S to S' . Therefore we can also view Z as a presentation 2-complex for $(G_d)_A$ with generators S' . Let \tilde{Z} be the cover of Z with fundamental group $(G_1)_A$. Since Δ' is a finite graph, \tilde{Z} is a finite-sheeted cover and is therefore a finite

complex. Then we use the Seifert–Van Kampen theorem to write down a finite presentation for the fundamental group of \tilde{Z} , which is $(G_1)_A$. Since every step in this construction can be done effectively, this is an algorithm to produce a finite presentation. \square

Example 5.7. We consider a concrete example to illustrate the above algorithms. It also happens that this is an example where the Schreier graph we describe is disconnected, and where a pair of matrices are in the same G_d -orbit but in different G_1 -orbits. Consider the following matrices:

$$A = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}, \quad N = \begin{pmatrix} 0 \\ 0 \\ 2 \end{pmatrix}, \quad Q = \begin{pmatrix} 1 & 0 & -\frac{1}{2} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then $A, N \in M_{3,1}(\mathbb{Z})$ and $Q \in G_2 = \mathrm{GL}(2, \mathbb{Z}) \ltimes M_{2,1}(\frac{1}{2}\mathbb{Z})$. We see that A is not in $G_{\mathbb{Q}}$ -normal form, but N is, and $N = QA$. The stabilizer of N in G_2 is a copy of $\mathrm{GL}(2, \mathbb{Z})$ generated by $S = \{a, b, c\}$ with

$$a = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad b = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad c = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Using the fact that $\mathrm{SL}(2, \mathbb{Z})$ is the amalgamated free product $(\mathbb{Z}/4\mathbb{Z}) *_{\mathbb{Z}/2\mathbb{Z}} (\mathbb{Z}/6\mathbb{Z})$ (see for example, Serre [12, page 35]) with the elements of order 4 and 6 given by $ba^{-1}b$ and $a^{-1}b$ respectively, it is straightforward to derive a presentation for $\mathrm{GL}(2, \mathbb{Z}) = (G_2)_N$ as follows:

$$\langle a, b, c \mid c^2 = 1, (a^{-1}b)^3 = (ba^{-1}b)^2, (a^{-1}b)^6 = 1, cac = a^{-1}, cbc = b^{-1} \rangle.$$

As usual, $G_1 = \mathrm{GL}(2, \mathbb{Z}) \ltimes M_{2,1}(\mathbb{Z})$. The Schreier graph Δ of G_1 in G_2 with respect to S is displayed in Figure 1. By inspecting Δ , we learn that the fundamental group of the connected component of $Q^{-1}G_1$, based at $Q^{-1}G_1$, is generated by the following seven elements:

$$\{a, c, b^2, bcb^{-1}, ba^2b^{-1}, bac a^{-1}b^{-1}, baba^{-1}b^{-1}\}.$$

Therefore the stabilizer $(G_1)_A$ is generated by the conjugates of these seven elements by Q^{-1} :

$$\left\{ a, c, \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & -1 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} -1 & 0 & 1 \\ -2 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} -1 & 2 & 1 \\ -2 & 3 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \right. \\ \left. \begin{pmatrix} -3 & 2 & 2 \\ -4 & 3 & 2 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 3 & -1 & -1 \\ 4 & -1 & -2 \\ 0 & 0 & 1 \end{pmatrix} \right\}.$$




Figure 1: The Schreier graph of $GL(2, \mathbb{Z}) \ltimes M_{2,1}(\mathbb{Z})$ in $GL(2, \mathbb{Z}) \ltimes M_{2,1}(\frac{1}{2}\mathbb{Z})$ with respect to $\{a, b, c\}$. Each vertex is labeled by the image of its coset under ρ . The vertex $Q^{-1}G_1$ is labeled $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

There is a presentation 2-complex for $(G_1)_A$ that has the connected component of QG_1 in Δ as its 1-skeleton, and which is a three-sheeted cover of the presentation 2-complex for $(G_2)_N$ corresponding to the presentation given above. In particular, $(G_1)_A$ has a finite presentation with seven generators and fifteen relators.

There is another interesting observation we can make by looking at Δ . Even though A and N are both in $M_{3,1}(\mathbb{Z})$ and A and N are in the same orbit under G_2 , the vertices G_1 and QG_1 are in different connected components of Δ and therefore A and N are in different orbits under the action of G_1 . This shows that it is not enough simply to check whether matrices are in the same $G_{\mathbb{Q}}$ -orbit.

6 Peak reduction

6.1 Preliminary notions

Throughout this section we use the generalized Whitehead automorphisms Ω defined in the introduction. The definition of peak here does not exactly match the definition used in Day [3], but the current definition is more symmetric.

Definition 6.1. A *peak* is a triple (W, α, β) , where W is a tuple of conjugacy

classes and α and β are automorphisms of A_Γ and

$$|\alpha \cdot W| \leq |W| \quad \text{and} \quad |\beta \cdot W| \leq |W|,$$

with at least one of these inequalities being strict. The *height* of a peak is $|W|$. In this paper, the automorphisms α and β in a peak are assumed to be in Ω .

A *lowering* of a peak (W, α, β) is a factorization

$$\beta\alpha^{-1} = \gamma_k \dots \gamma_1$$

of $\beta\alpha^{-1}$ by automorphisms $\gamma_1, \dots, \gamma_k$, such that all the lengths of the intermediate images of $\alpha \cdot W$ are strictly lower than that of W :

$$|\gamma_i \dots \gamma_1 \alpha \cdot W| < |W|$$

for $i = 1, \dots, k-1$. In this paper, the automorphisms in a lowering factorization of a peak are always elements of Ω .

The goal of this section is to prove the following lemma:

Main Lemma 6.2. *Suppose (W, α, β) is a peak and α and β are generalized Whitehead automorphisms. Then this peak can be lowered using a factorization of $\beta\alpha^{-1}$ by generalized Whitehead automorphisms.*

Theorem 1.3 follows immediately from this lemma:

Proof of Theorem 1.3 from Main Lemma 6.2. This is a standard argument; it appears in detail as the “Proof of part (3) of Theorem B” on pages 836–837 of [3]. In short, we express α as a product of generalized Whitehead automorphisms (for example, as a product of Laurence–Servatius generators using Laurence’s theorem) and then repeatedly alter the factorization using Main Lemma 6.2. Whenever we see a peak of maximal height, we replace it with a peak-lowering factorization. Each application of the Main Lemma reduces either the number of peaks of maximal height or reduces the maximum of the heights of the peaks in the factorization. The procedure will terminate with a factorization of α as product of automorphisms from Ω such that no peaks appear in the sequence of lengths of the intermediate images of W . Such a factorization satisfies the conclusion of the theorem. \square

In our algorithm to lower peaks, we break into cases depending on the properties of the automorphisms α and β in the given peak (W, α, β) . We often consider the support of automorphisms, defined in Definition 3.10. Using this, we can explain the structure of a generalized Whitehead automorphism.

Lemma 6.3. *Suppose $\alpha \in \Omega_{[a]}$ for some vertex a . Then α is a product of dominated transvections and partial conjugations with multipliers in $[a]$ and inversions in elements of $[a]$. In particular:*

- *if b is adjacent to a with $[b] \neq [a]$ and α does not fix b , then a dominates b .*
- *if Y is a connected component of $\Gamma \setminus \text{st}(a)$ with at least two elements, then either $Y^{\pm 1} \cap \text{supp}(\alpha) = \emptyset$ or $Y^{\pm 1} \subset \text{supp}(\alpha)$.*

Proof. We know that $\Omega_{[a]}$ is isomorphic to a semidirect product $\text{GL}(k, \mathbb{Z}) \ltimes M_{k,l}(\mathbb{Z})$, where the short-range dominated transvections between elements of $[a]$ and inversions in $[a]$ generate the $\text{GL}(k, \mathbb{Z})$ -factor and the other dominated transvections and partial conjugations with multipliers in $[a]$ generate the $M_{k,l}(\mathbb{Z})$ -factor. This $M_{k,l}(\mathbb{Z})$ is a free abelian group. In particular, we can express α as $\alpha_2 \alpha_1$, where α_1 acts trivially on $\Gamma \setminus [a]$ (it is in the $\text{GL}(k, \mathbb{Z})$ -factor) and α_2 acts trivially on $[a]$ (it is in the $M_{k,l}(\mathbb{Z})$ -factor).

If b is adjacent to a with $[b] \neq [a]$, then the only possible generators of $\Omega_{[a]}$ that can alter b are transvections, if they exist. So if α does not fix b , there is a dominated transvection with multiplier a acting on b , and therefore a dominates b .

For Y a connected component of $\Gamma \setminus \text{st}(a)$, the only possible generators of $\Omega_{[a]}$ that do not leave Y pointwise fixed are the partial conjugations that conjugate Y by an element of $[a]$. If one of these appears in a minimal-length factorization of α_2 , then $Y^{\pm 1} \subset \text{supp}(\alpha)$. If none do, then $Y^{\pm 1} \cap \text{supp}(\alpha) = \emptyset$. \square

We need some basic observations about the interactions of Whitehead automorphisms, which we gather together in the following lemma.

Lemma 6.4. *Suppose a and b are vertices of Γ and $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$. Then:*

- *If a is adjacent to b and $\alpha|_{[b]}$ is not the identity and $\beta|_{[a]}$ is not the identity, then $[a] = [b]$.*
- *If a is not adjacent to b and α fixes each element of $[b]$ then α fixes each element of $\text{st}(b)$.*
- *If a is not adjacent to b then α and β both fix all of $\text{st}(a) \cap \text{st}(b)$.*
- *If a is not adjacent to b and a dominates b , then β is a long-range automorphism.*

- If a is not adjacent to b and $\alpha|_{[b]}$ is not the identity and $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ then β is a long-range automorphism.

Proof. For the first item, suppose for contradiction that $[a] \neq [b]$. We factor α and β as products of Laurence generators. Since $\alpha|_{[b]}$ is not the identity, some factor of α is a dominated transvection that replaces some $c \in [b]$ with cd for some $d \in [a]$. In particular, this means that b dominates a . Reversing the roles of α and β , we see that a dominates b as well. Then by definition $[a] = [b]$, a contradiction.

For the second item, we write $\alpha = \alpha_1 \alpha_2$ where α_1 is a product of partial conjugations and α_2 is a product of transvections and of inversions of elements of $[a]$. Let $x \in \text{st}(b) \setminus [b]$. If α_1 does not fix x , then x is not adjacent to a and α_1 conjugates the entire connected component of x in $\Gamma \setminus \text{st}(a)$ by the same nontrivial element of $\langle [a] \rangle$. Necessarily b is in the same connected component and is therefore not fixed by α_1 . But since x is adjacent to b but not to a , we know that a does not dominate b , and therefore α_2 fixes b . This implies that α does not fix b , a contradiction. If α_2 does not fix x , then there is a transvection multiplying an element of $[a]$ onto x . Then a dominates x . But since x is adjacent to b , this implies a is adjacent to b , a contradiction. So neither α_1 nor α_2 can alter x , which was arbitrary, and therefore α fixes $\text{st}(b)$.

The third item is similar. If α does not fix an element x of $\text{st}(a) \cap \text{st}(b)$, then there must be a transvection multiplying an element of $[a]$ onto x , so that a dominates x . But if a dominates x and x is adjacent to b , then a is adjacent to b , a contradiction. Similarly β must fix every element of $\text{st}(a) \cap \text{st}(b)$.

To show the fourth item, we suppose that a dominates b . If x is adjacent to b and β does not fix x , then there is a transvection factor of β that multiplies an element of $[b]$ onto x . Then b dominates x . Since a dominates b , this implies that a dominates x , and therefore that a is adjacent to b , because x is. This is a contradiction, so therefore β fixes every vertex adjacent to b .

Now we explain the fifth item. If α does not fix b , then either a dominates b , or the component Y of a in $\Gamma \setminus \text{st}(b)$ has at least two elements and α conjugates every element of Y by the same nontrivial word in $[a]$. If a dominates b , then we have already shown that β must be long-range. So we suppose Y has at least two elements. Suppose x is adjacent to b . If β does not fix x , then b dominates x (as explained in the previous item). If x were adjacent to a , then b would be adjacent to a , which it is not. So x is not adjacent to a , and therefore x is an element of Y . However, if x is an element of Y , then x is in $\text{supp}(\beta)$, a contradiction since $\text{supp}(\beta) \cap \text{supp}(\alpha) = \emptyset$ and

$x \in \text{supp}(\alpha)$. So β fixes each vertex adjacent to b . \square

The following observation is easy and needs no proof. It greatly reduces the cases we need to consider.

Lemma 6.5. *If (W, α, β) is a peak, then so is (W, β, α) . If one of these peaks can be lowered, then so can the other.*

Now we outline the proof of Main Lemma 6.2, which is broken into lemmas filling out the rest of this section.

Proof of Main Lemma 6.2. Since (W, α, β) is a peak, at least one of α or β shortens W . Since permutation automorphisms do not shorten W , at least one of α or β is not a permutation automorphism. Since we may swap α and β , we assume that α is not a permutation, so $\alpha \in \Omega_{[a]}$ for some vertex a of Γ . If β is a permutation automorphism, then conjugate α by β to lower the peak as described in Lemma 6.10.

So we assume that neither α nor β is a permutation automorphism; then $\beta \in \Omega_{[b]}$ for some vertex b in Γ . If $[a] = [b]$, then we lower the peak by replacing the length-two factorization $\beta\alpha^{-1}$ with the length-one factorization given by its product, as explained in Lemma 6.11.

Next we suppose that a and b are adjacent to each other in Γ , but that $[a] \neq [b]$. By Lemma 6.4, this implies that $\alpha|_{[b]}$ is the identity or $\beta|_{[a]}$ is the identity. We suppose that $\alpha|_{[b]}$ is the identity; then Lemma 6.12 shows how to lower the peak. This uses a Steinberg relation explained in Lemma 4.7.

So we assume that a and b are not adjacent in Γ . If a dominates b and b dominates a , then by Lemma 6.4, both α and β are long-range automorphisms. In this case, we can lower the peak (in fact, fully reduce it) using Theorem 4.6. So we suppose that b dominates a but a does not dominate b . Again by Lemma 6.4, α is a long-range automorphism. Proposition 6.17 explains how to lower such a peak. Finally, we suppose that b does not dominate a and a does not dominate b . Then Proposition 6.15 shows that the peak can be lowered. \square

6.2 Change of length under generalized Whitehead automorphisms

To show that the algorithm for lowering peaks works as desired, we need to show that certain factorizations really are peak lowering. To do this we need a good understanding of the effect that generalized Whitehead automorphisms have on lengths of tuples of cyclic words. Next we prove some results about this.

We recall the action of $\Omega_{[a]}$ on syllables with respect to a given in Definition 3.6.

Lemma 6.6. *Suppose W is a tuple of cyclic words and T is a syllable decomposition of W with respect to a and $\alpha \in \Omega_{[a]}$. Then*

$$|\alpha \cdot W| - |W| = |\alpha \cdot T| - |T|.$$

In particular, we can compute the difference $|\alpha \cdot W| - |W|$ by computing the differences $|\alpha \cdot t| - |t|$ over all syllables t of T and summing.

Proof. This is an easy corollary of Lemma 3.8; the details are left to the reader. \square

In the following, when we talk about an endpoint of a syllable *cud* being in the support of an automorphism, we are only talking about whether c and d^{-1} are in the support, but not about c^{-1} or d . This is because the action of automorphisms on syllables has nothing to do with the action on the far sides of the endpoints. The following proposition is used to show that factorizations coming from the relations in Lemma 4.7 can lower peaks.

Proposition 6.7. *Let a and b be vertices of Γ with $[a] \neq [b]$. Suppose $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$ satisfy the hypotheses of Lemma 4.7 (α fixes $[b]$, and either a is adjacent to b , or $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ and β fixes $[a]$). Let $\gamma = \alpha\beta\alpha^{-1}$, which is in $\Omega_{[b]}$ by Lemma 4.7. Suppose W is a tuple of conjugacy classes. Then*

$$|W| - |\beta \cdot W| = |\alpha \cdot W| - |\alpha\beta \cdot W|.$$

Further, if (W, α, β) is a peak, then $|\gamma\alpha \cdot W| < |W|$.

Since $\gamma\alpha = \alpha\beta$, we can also express $\alpha\beta \cdot W$ as $\gamma\alpha \cdot W$ in the equation above.

Proof of Proposition 6.7. Let $T = (t_1, \dots, t_N)$ be a syllable decomposition of W with respect to $[b]$. We form a syllable decomposition T' of $\alpha \cdot W$ with respect to $[b]$ by the following process: first we form the representative of W associated to T , then we apply α to this representative by inserting elements from $[a]$ where appropriate, then we delete inverse pairs of elements of $[a]$ that commute with all intervening letters, and then we split the resulting representative into syllables. We note that no collapsing can happen in this process: if $c \notin [a]$ cancels with its inverse after we apply α , then it could have canceled before applying α , thus contradicting the stipulation from

the definition of syllable decomposition that the associated representative be graphically reduced. We prove the proposition by defining a relation f between the syllables in T and the syllables T' . This f will be an injective partial function.

Consider syllable t_i from T . Either t_i is a linear syllable $t_i = c_i v_i u_i d_i$ or a cyclic syllable $t_i = v_i u_i$, with $v_i \in \langle [b] \rangle$, $u_i \in \langle \text{st}(b) \setminus [b] \rangle$ and $c_i, d_i \in (\Gamma \setminus \text{st}(b))^{\pm 1}$. We consider what happens to t_i in passing from W to $\alpha \cdot W$. The hypotheses of Lemma 4.7 give us two cases for α and β . In the first case, a is adjacent to b , so acting by α cannot alter syllable boundaries, and we define f to be the map that sends each syllable in T to the corresponding syllable with the same boundaries in T' .

In the other case, a is not adjacent to b and $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$. So by Lemma 6.4, α fixes $v_i u_i$. In this case, $\beta|_{[a]}$ is the identity. Also, if c_i or d_i^{-1} is in $\text{supp}(\beta)$, it is not in $[a]$ and is not in $\text{supp}(\alpha)$, so action by α cannot cancel it away or insert an element of $[a]$ between it and $v_i u_i$. So in the second case, if c_i or d_i^{-1} or any element of u_i is in $\text{supp}(\beta)$ then all of these land in the same syllable of T' ; we declare f to send t_i to this syllable.

In the second case, if c_i and d_i^{-1} are not in $\text{supp}(\beta)$ and u_i does not contain any elements of $\text{supp}(\beta)$, then we leave t_i out of the domain of f . This completes the definition of f . We note that f is injective. This is obvious in the first case; in the second case, if two or more syllables merge to form the syllable t'_j , then at most one of them can be in the domain of f . This is because $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ and because elements of $\text{st}(a) \cap \text{st}(b)$ cannot be in $\text{supp}(\beta)$: to break a syllable boundary, we must have the far endpoint being in $\text{supp}(\alpha)$ and everything in $\text{st}(b)$ on one side of the boundary in $\text{st}(b) \cap \text{st}(a)$.

Claim. *If t_i is in the domain of f , then $|t_i| - |\beta \cdot t_i| = |f(t_i)| - |\gamma \cdot f(t_i)|$.*

First we consider the case where a is adjacent to b . Then $f(t_i)$ differs from t_i in that some elements of $[a]$ have been added and deleted from u_i . We use u'_i to denote the part of $f(t_i)$ that is a word from $\text{st}(b) \setminus [b]$. The syllable $\beta \cdot t_i$ differs from t_i in that elements of $[b]$ have been added and deleted from v_i . Let v'_i denote the part of $\beta \cdot t_i$ that is a word from $[b]$. Since $\gamma\alpha = \alpha\beta$, the syllable $\gamma \cdot f(t_i)$ can be computed by applying α to the syllable $\beta \cdot t_i$ (in the same manner that we obtained $f(t_i)$ by applying α to t_i). Since the elements of $\text{supp}(\alpha)$ in t_i are in c_i , d_i^{-1} and u_i (since α fixes $[b]$), these are unchanged in $\beta \cdot t_i$, and therefore $\gamma \cdot f(t_i)$ differs from $\beta \cdot t_i$ in that u_i is replaced by u'_i (the same u'_i as above). So the $\text{st}(b)$ -parts of t_i , $\beta \cdot t_i$, $f(t_i)$ and $\gamma \cdot f(t_i)$ are $u_i v_i$, $u_i v'_i$, $u'_i v_i$ and $u'_i v'_i$ respectively. Since these are reduced products, this proves the claim in this case.

Next we consider the case where a is not adjacent to b , the supports of α and β are disjoint, and β fixes $[a]$. In this case, $\gamma = \beta$. The elements of $f(t_i)$ in $\text{supp}(\beta)$ are exactly the same as those in t_i , since, as noted above, any elements merged in from other syllables cannot contain anything from $\text{supp}(\beta)$. Similarly, parts merged in from other syllables cannot contain elements of $[b]$. By Lemma 6.4, α fixes $v_i u_i$. Of course, $\beta \cdot t_i$ differs from t_i in that v_i is replaced by v'_i , where v'_i is determined by v_i and the elements of t_i in $\text{supp}(\beta)$. Since $f(t_i)$ contains the same $[b]$ -part v_i and the same elements of $\text{supp}(\beta)$, it follows that $\beta \cdot t_i$ will be the same as $f(t_i)$, but with v_i replaced by the same v'_i just mentioned. In particular, the claim follows in this case.

Claim. *If t_i is not in the domain of f , then $t_i = \beta \cdot t_i$, and if t'_j is not in the range of f , then $t'_j = \gamma \cdot t'_j$.*

In the first case, f is a bijective total function and there is nothing to show. In the second case, the statement about the domain of f is obvious: if t_i is not in the domain, then it contains nothing from $\text{supp}(\beta)$, and therefore β does nothing to it. In the second case, if t'_j is some syllable in T' not in the range of f , then no part of it is in $\text{supp}(\beta)$. This is because α did not create new elements of $\text{supp}(\beta)$ in acting on W , and if some t_i has some elements of $\text{supp}(\beta)$, all these elements of $\text{supp}(\beta)$ end up in its correspondent under f . The action of γ is to place cancelling copies of v and v^{-1} in the syllable, proving the claim.

Now we finish the proof of the proposition. Using Lemma 6.6, we have shown that

$$|W| - |\beta \cdot W| = |\alpha \cdot W| - |\gamma \alpha \cdot W|,$$

since we have matched up all the syllables that change length on one side with the syllables that change length on the other side, and shown that they change length by the same amount. If we suppose that (W, α, β) is a peak, then by definition, we know that $2|W| < |\beta \cdot W| + |\alpha \cdot W|$ (this is from summing the two inequalities in the definition of a peak and using the stipulation that one of them is strict). Combining these two statements, we obtain $|W| > |\gamma \alpha \cdot W|$. \square

For the next proposition, we recall the technique from Day [3] for computing the change in length of classic long-range Whitehead automorphisms. We defined classic Whitehead automorphisms in Section 4.3. Fix a tuple of conjugacy classes W . For c in a vertex of Γ and A and B subsets of $X^{\pm 1}$, we use the notation $\langle A, B \rangle_{W,c}$ to denote the number of instances of subwords of the forms due^{-1} or eud^{-1} in a graphically reduced representative for W ,

where $d \in A \setminus \text{st}(c)$, $e \in B \setminus \text{st}(c)$, and u is a word in $\text{st}(c)$. This number is nonnegative and does not depend on the choices involved. As a function of A and B , it is additive over disjoint sets in both inputs. Generally we will work with a specific W and suppress W from the notation. Let γ be a long-range classic Whitehead automorphism with multiplier c and write $C = \text{supp}(\gamma) + c$. We write c for $\{c\}$, “-” for differences of sets, “+” for disjoint unions and C' for $X^{\pm 1} \setminus C$ in the following computations. This bracket has the following useful property (see Day [3, Lemma 3.16, Lemma 3.17]):

$$|W| - |\gamma \cdot W| = \langle c, C - c \rangle_c - \langle C', C - c \rangle_c.$$

Since $\langle c, c \rangle_c = 0$ (W is graphically reduced), we can rewrite this as

$$|W| - |\gamma \cdot W| = \langle c, X^{\pm 1} \rangle_c - \langle C', C \rangle_c.$$

We use this bracket to detect sets of vertices for constructing Whitehead automorphisms for factorizations used in the peak reduction algorithm. We need a couple of lemmas for this purpose.

Proposition 6.8. *Suppose α and β are classic long-range Whitehead automorphisms with multipliers a and b respectively. Suppose that a is not adjacent to b and b dominates a . Finally, we suppose that $\alpha(b) = b$ and a is in $\text{supp}(\beta)$. Let W be a tuple of conjugacy classes. Let A denote $\text{supp}(\alpha) + a$ and let B denote $\text{supp}(\beta) + b$. Then:*

- *there is a classic long-range Whitehead automorphism β_1 with multiplier b^{-1} and support $\text{supp}(\beta_1) = A' \cap B' - b^{-1}$;*
- *there is a classic long-range Whitehead automorphism α_1 with multiplier a and support $\text{supp}(\alpha_1) = A \cap B - a$; and*
- *we have the following inequality for the changes in length of W under these automorphisms:*

$$(|W| - |\beta_1 \cdot W|) + (|W| - |\alpha_1 \cdot W|) \geq (|W| - |\beta \cdot W|) + (|W| - |\alpha \cdot W|).$$

Proof. First we explain why the automorphisms β_1 and α_1 exist. Since b dominates a , b dominates every vertex that a dominates, and if Y is a connected component of $\Gamma \setminus \text{st}(a)$, Y is a union of vertices adjacent to b , vertices dominated by b , and connected components of $\Gamma \setminus \text{st}(b)$. Therefore β_1 can be expressed as a product of Laurence generators and is a well defined automorphism.

If a dominates b , then a and b dominate the same vertices non-adjacently and $\Gamma \setminus \text{st}(a)$ and $\Gamma \setminus \text{st}(b)$ have the same connected components with two or more vertices. Then $A \cap B - a$ is a union of vertices that a dominates non-adjacently and components of $\Gamma \setminus \text{st}(a)$ with two or more vertices, and therefore α_1 can be expressed as a product of Laurence generators with multiplier a and is well defined. Now suppose that a does not dominate b . Then since α fixes b , α must fix the entire connected component of b in $\Gamma \setminus \text{st}(a)$. If c is a vertex in $A \cap B$ and a does not dominate c , then the entire connected component Y of c in $\Gamma \setminus \text{st}(a)$ is in A . If b dominates c , then there is a path from c to b outside of $\text{st}(a)$, and therefore b and c are both in Y . However, this is a contradiction, since α must conjugate every element of Y by a . So b does not dominate c , which means that the entire connected component Z of c in $\Gamma \setminus \text{st}(b)$ is in B . Since $Y \subset Z$ (since b dominates a), this means that $Y \subset A \cap B$. In particular, α_1 can be expressed as a product of Laurence generators and is well defined.

Now we consider lengths of images of W . From the comments preceding the lemma, we know

$$|W| - |\alpha \cdot W| = \langle a, X^{\pm 1} \rangle_a - \langle A', A \rangle_a,$$

$$|W| - |\beta \cdot W| = \langle b, X^{\pm 1} \rangle_b - \langle B', B \rangle_b,$$

and

$$|W| - |\beta_1 \cdot W| = \langle b^{-1}, X^{\pm 1} \rangle_b - \langle A' \cap B', (A' \cap B')' \rangle_b.$$

We expand:

$$\begin{aligned} \langle A' \cap B', (A' \cap B')' \rangle_b &= \langle A' \cap B', A \cap B \rangle_b + \langle A' \cap B', A' \cap B \rangle_b \\ &\quad + \langle A' \cap B', A \cap B' \rangle_b \end{aligned}$$

and:

$$\begin{aligned} \langle B', B \rangle_b &= \langle A \cap B', A \cap B \rangle_b + \langle A \cap B', A' \cap B \rangle_b + \langle A' \cap B', A \cap B \rangle_b \\ &\quad + \langle A' \cap B', A' \cap B \rangle_b. \end{aligned}$$

In particular, we see

$$\begin{aligned} (|W| - |\beta_1 \cdot W|) - (|W| - |\beta \cdot W|) &= \langle A \cap B, A \cap B' \rangle_b - \langle A' \cap B', A \cap B' \rangle_b \\ &\quad + \langle A \cap B', A' \cap B \rangle_b. \end{aligned}$$

Let C_W be the number of subwords of our graphically reduced representative for W of the form auc^{-1} , where c is in $A - a$ and u is in $\langle \text{st}(b) \rangle$ but not in $\langle \text{st}(a) \rangle$ (so u contains some letter not commuting with a). Then

$$\langle A \cap B, A \cap B' \rangle_b = \langle A \cap B, A \cap B' \rangle_a + C_W.$$

On the other hand, if cud^{-1} is counted by $\langle A' \cap B', A \cap B' \rangle_b$, meaning that $c \in A' \cap B'$, $u \in \langle \text{st}(b) \rangle$ and $d \in A - a$, then either $u \in \langle \text{st}(a) \rangle$ and cud^{-1} is counted by $\langle A' \cap B', A \cap B' \rangle_a$, or else $u = u_1 eu_2$ with $u_1 \in \langle \text{st}(a) \rangle$ and $e \in \text{st}(b) \setminus \text{st}(a)$. In the latter case, $cu_1 e$ is counted by $\langle A' \cap B', A \cap B' \rangle_a$, since e is in $A' \cap B'$ (e is fixed by β because β is long range and by α because otherwise α would not fix b). So each subword counted by $\langle A' \cap B', A \cap B' \rangle_b$ is counted exactly once by $\langle A' \cap B', A \cap B' \rangle_a$, so

$$\langle A' \cap B', A \cap B' \rangle_b = \langle A' \cap B', A \cap B' \rangle_a.$$

Then

$$(|W| - |\beta_1 \cdot W|) - (|W| - |\beta \cdot W|) = \langle A \cap B, A \cap B' \rangle_a - \langle A' \cap B', A \cap B' \rangle_a + C_W + \langle A \cap B', A' \cap B \rangle_b.$$

By an argument parallel to the one for β_1 , we deduce

$$(|W| - |\alpha_1 \cdot W|) - (|W| - |\alpha \cdot W|) = \langle A \cap B', A' \cap B' \rangle_a - \langle A \cap B, A \cap B' \rangle_a + \langle A' \cap B, A \cap B' \rangle_a.$$

Summing the two equations, we see that

$$(|W| - |\beta_1 \cdot W|) - (|W| - |\beta \cdot W|) + (|W| - |\alpha_1 \cdot W|) - (|W| - |\alpha \cdot W|)$$

is

$$C_W + \langle A \cap B', A' \cap B \rangle_b + \langle A \cap B', A' \cap B \rangle_a,$$

a nonnegative number. \square

In the following, when we write $|\alpha|_O$, we mean the minimum length of a representative of the class of α in $\text{Out}(A_\Gamma)$ as a product of Laurence generators.

Corollary 6.9. *Suppose (W, α, β) is a peak with $\alpha \in \Omega_{[a]}$, $\beta \in \Omega_{[b]}$ classic long-range Whitehead automorphisms, with a not adjacent to b , and b dominating a , but a not dominating b . Suppose $a \in \text{supp}(\beta)$ and $\alpha(b) = b$. Then either*

- *there is a classic long-range Whitehead automorphism β_1 with multiplier b^{-1} , with $|\beta_1^{-1}\beta|_O < |\beta|_O$ as a product of Laurence generators, with $\beta_1(a)$ equal to ba or a , and with $|\beta_1 \cdot W| < |W|$, or*
- *there is a classic long-range Whitehead automorphism α_1 with multiplier a , with $\text{supp}(\alpha_1) \subset \text{supp}(\beta)$, and with $|\alpha_1 \cdot W| < |W|$.*

Further, if $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$, then the first option holds.

Proof. We apply Proposition 6.8. Using the definition of a peak,

$$(|W| - |\beta_1 \cdot W|) + (|W| - |\alpha_1 \cdot W|) \geq (|W| - |\beta \cdot W|) + (|W| - |\alpha \cdot W|) > 0.$$

So one of α_1 or β_1 strictly shortens W . Since α_1 and β_1 satisfy the other properties in the lemma, this proves the main statement of the lemma. If $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$, then α_1 is defined to be the trivial automorphism and cannot shorten W . Therefore in this case, it is β_1 that shortens W . \square

6.3 The algorithm: basic cases

Now we begin considering cases for lowering peaks. For the rest of this section, we suppose we have peak (W, α, β) with α and β generalized Whitehead automorphisms.

Lemma 6.10. *If α is a permutation automorphism, then the peak (W, α, β) can be lowered.*

Proof. Since α is a permutation automorphism, $\alpha \cdot W$ is the same length as W . Then the definition of a peak demands that $|\beta \cdot W| < |W|$, and therefore $\beta \in \Omega_{[b]}$ for some vertex b of Γ . Then $\beta\alpha^{-1} = \alpha^{-1} \cdot \alpha\beta\alpha^{-1}$ is a peak-lowering factorization: $\alpha\beta\alpha^{-1}$ is in $\Omega_{[\alpha(b)]}$ and $|\alpha\beta \cdot W| = |\beta \cdot W| < |W|$ (again because α is a permutation automorphism). \square

Lemma 6.11. *If α and β are both in $\Omega_{[a]}$ for a vertex of Γ , then (W, α, β) can be lowered.*

Proof. Set $\gamma = \beta\alpha^{-1}$. Then $\beta\alpha^{-1} = \gamma$ is a peak-lowering factorization: $\gamma \in \Omega_{[a]}$ and the condition on intermediate lengths of images of W is vacuous for factorizations of length one. \square

For the rest of this section we assume that (W, α, β) is a peak with $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$ for distinct vertices a and b in Γ .

Lemma 6.12. *Suppose a adjacent is adjacent to b and $\alpha_{[b]}$ is the identity. Then the peak (W, α, β) can be lowered.*

Proof. By Lemma 4.7, we know that the element $\gamma = \alpha\beta\alpha^{-1}$ is in $\Omega_{[b]}$. We use the factorization $\beta\alpha^{-1} = \alpha^{-1}\gamma$. To show that this is peak-lowering, it is enough to show that $|\gamma\alpha \cdot W| < |W|$. However, we have already done this in Proposition 6.7 \square

Lemma 6.13. *Suppose a is not adjacent to b , $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ and $\alpha|_{[b]}$ and $\beta|_{[a]}$ are both identity maps. Then the peak (W, α, β) can be lowered.*

Proof. Again we use $\gamma = \alpha\beta\alpha^{-1}$, which is in $\Omega_{[b]}$ by Lemma 4.7. Our hypotheses imply that α and β commute, so that $\gamma = \beta$. Our peak-lowering factorization is

$$\beta\alpha^{-1} = \alpha^{-1}\beta.$$

To show this, of course we need to show that $|\beta\alpha \cdot W| < |W|$. We have already shown this in Proposition 6.7. \square

6.4 General cases with non-adjacent multipliers

Lemma 6.14. *Suppose (W, α, β) is a peak with $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega$, and $\gamma \in \Omega_{[a]}$ is an inner automorphism of A_Γ . Then $\gamma\alpha$ and $\alpha\gamma$ are both in $\Omega_{[a]}$ and $(W, \gamma\alpha, \beta)$ and $(W, \alpha\gamma, \beta)$ are both peaks and can be lowered if and only if (W, α, β) can be lowered.*

Proof. The claim that $\gamma\alpha$ and $\alpha\gamma$ are in $\Omega_{[a]}$ is obvious. We note that $\gamma\alpha = \alpha\gamma'$ for a possibly different inner automorphism $\gamma' \in \Omega_{[a]}$, so it is enough to prove the lemma for $(W, \gamma\alpha, \beta)$. Also, it is enough to show only the “if” direction of the statement.

First of all, if (W, α, β) is a peak, then $(W, \gamma\alpha, \beta)$ is as well, since action by inner automorphisms has no effect on cyclic words. Now we suppose that $\delta_1, \dots, \delta_k \in \Omega$ and

$$\beta\alpha^{-1} = \delta_k \cdots \delta_1$$

is a lowering of (W, α, β) . Let l be one of $1, \dots, k$ such that $\delta_l \cdots \delta_1 \alpha \cdot W$ is minimal length among all $\delta_i \cdots \delta_1 \alpha \cdot W$. Since a conjugate of an inner automorphism by another automorphism is still an inner automorphism, and since every inner automorphism is a product of inner automorphisms from Ω , we can find inner automorphisms $\gamma_1, \dots, \gamma_m$ in Ω such that

$$\beta\alpha^{-1}\gamma^{-1} = \delta_k \cdots \delta_{l+1}\gamma_m \cdots \gamma_1\delta_l \cdots \delta_1.$$

Since inner automorphisms do not change the length of cyclic words, this is a lowering of the peak $(W, \gamma\alpha, \beta)$. \square

If α is a classic long-range Whitehead automorphism with multiplier a , we will sometimes consider the *complement* of α . This is the classic long-range Whitehead automorphism with multiplier a , with the property that the union of the support of α and the support of its complement is $(\Gamma \setminus \text{st}(a))^{\pm 1}$. It follows that the product of α and its complement is inner,

so Lemma 6.14 implies that we may freely replace an automorphism in a peak with its complement.

Proposition 6.15. *Suppose a is not adjacent to b , b does not dominate a and a does not dominate b . Then the peak (W, α, β) can be lowered.*

Proof. We prove the proposition by proving it in successively more general cases, with each case building on the last. First we prove:

Case. *The proposition is true if β is a classic long-range Whitehead automorphism with multiplier b , $\beta|_{[a]}$ is the identity, $\alpha|_{[b]}$ is not the identity and $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$.*

Since a does not dominate b , there is some nontrivial element $v \in \langle [a] \rangle$ such that $\alpha(b) = vbv^{-1}$. We define a new automorphism $\alpha_1 \in \Omega_{[a]}$ by $\alpha_1(c) = v^{-1}\alpha(c)v$ for every $c \in \Gamma$. Then α_1 and α differ by an inner automorphism, and to prove the claim, it is enough to show that the peak (W, α_1, β) can be lowered.

Claim. *If $c^{\pm 1} \in \text{supp}(\beta)$, then either a dominates c or $Y^{\pm 1} \subset \text{supp}(\beta)$, where Y is the connected component of c in $\Gamma \setminus \text{st}(a)$.*

Suppose that c is a vertex of Γ with c or $c^{-1} \in \text{supp}(\beta)$ and a does not dominate c . Let Y be the connected component of c in $\Gamma \setminus \text{st}(a)$. We want to show that Y is a subset of a connected component of $\Gamma \setminus \text{st}(b)$ with at least two vertices. This is enough, since by Lemma 6.3 this implies that $Y^{\pm 1} \subset \text{supp}(\beta)$. Then it is enough to show that b does not dominate c and that Y does not contain an element of $\text{st}(b)$.

If b dominates c , then since a does not dominate c , either b is adjacent to c or there is a vertex adjacent to b and c but not adjacent to a . If Y contains an element of $\text{st}(b)$, then there is a path from c to b outside of $\text{st}(a)$. In any of these three cases, it follows that b and c are in the same component of $\Gamma \setminus \text{st}(a)$. However, this contradicts Lemma 6.3: α fixes c and c^{-1} since β does not and $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$, but α does not fix b . This proves the claim.

By the claim, $\text{supp}(\beta)$ is a union of elements $c^{\pm 1}$ where a dominates c , and subsets $Y^{\pm 1}$ where Y is an entire connected component of $\Gamma \setminus \text{st}(a)$. In particular, the following automorphisms are products of Laurence generators and are well defined. Let $\alpha_2(c)$ be in $\{c, cv^{-1}, vc, vcv^{-1}\}$ for each vertex $c \in \Gamma$, with $\text{supp}(\alpha_2) = \text{supp}(\beta)$, and let $\alpha_3 = \alpha_1\alpha_2^{-1}$.

Now our goal is to show that $|\alpha_3 \cdot W| < |W|$. If this is true, then we will have that (W, α_3, β) is a peak with $\beta|_{[a]}$ and $\alpha_3|_{[b]}$ being identity maps and $\text{supp}(\alpha_3) \cap \text{supp}(\beta) = \emptyset$. In particular, Lemma 6.13 applies to lower

(W, α_3, β) . Appending α_2^{-1} to this factorization will give us a peak-lowering factorization for $\beta\alpha^{-1}$.

In fact we show that

$$|\alpha_1 \cdot W| - |\alpha_2 \cdot W| \geq |v|(|W| - |\beta \cdot W|).$$

Let T be a syllable decomposition of W with respect to $[b]$. For each syllable t_i in T , either $\beta \cdot t_i$ is longer than t_i or shorter than t_i or the same length. Let L_1 be the sum over syllables t_i that β lengthens of $|\beta \cdot t_i| - |t_i|$, and let S_1 be the sum over syllables t_i that β shortens of $|t_i| - |\beta \cdot t_i|$. Then

$$|W| - |\beta \cdot W| = S_1 - L_1.$$

Let T' be a syllable decomposition of $\alpha_1 \cdot W$ with respect to $[a]$, and similarly let L_2 be the total increase in length of syllables α_2^{-1} lengthens and let S_2 be the total decrease in length of syllables that α_2^{-1} shortens. Then

$$|\alpha_1 \cdot W| - |\alpha_3 \cdot W| = S_2 - L_2.$$

We assume without loss of generality that $\alpha_1 \cdot T$ and T' have the same associated representative of $\alpha_1 \cdot W$.

Suppose t_i is a syllable in T that decreases in length under β . Then t_i contains b to a nonzero power and exactly one endpoint of t_i is in $\text{supp}(\beta)$. Suppose $t_i = cb^kud$ and where $u \in \langle \text{st}(b) \setminus \{b\} \rangle$ and without loss of generality assume c is in $\text{supp}(\beta)$. Then the initial part $cb^{\pm 1}$ of t_i is a syllable of W with respect to $[a]$. The image of this syllable under α_1 is a syllable of $\alpha_1 \cdot W$ with respect to $[a]$; it is of the form $c\alpha_1(b)^{\pm 1}$. Clearly $\alpha_3 \cdot cb^{\pm 1}$ (acting on the syllable) is $cv^{-1}\alpha_1(b)^{\pm 1}$. This is a syllable of $\alpha_1 \cdot W$ in T' that shortens by $|v|$ under the action of α_2^{-1} . In particular, summing over all syllables in T that decrease in length under β , we have

$$|v|S_1 \leq S_2.$$

Suppose t'_j is a syllable in T' that increases in length under α_2^{-1} . Suppose $t'_j = cwud$ where $u \in \langle \text{st}(a) \setminus [a] \rangle$ and $w \in \langle [a] \rangle$. We note that α_2 is like a classic Whitehead automorphism in that $\alpha_2(c)$ must be one of $\{c, cv^{-1}, vc, vcv^{-1}\}$. In particular, since $\alpha_2^{-1} \cdot t'_j$ is longer than t_j , we must have exactly one of c and d^{-1} in $\text{supp}(\alpha_2)$, and it must be the case $w \neq 1$ that the v or v^{-1} that is inserted by α_2^{-1} does not cancel away completely into w . We assume without loss of generality that c is in $\text{supp}(\alpha_2) = \text{supp}(\beta)$ but d^{-1} is not. Clearly t'_j increases in length by at most $|v|$ under α_2^{-1} . We consider $\alpha_1^{-1} \cdot t'_j$, which is a syllable of W with respect to $[a]$. Then

$\alpha_1^{-1} \cdot t'_j = cw'ud$, where $w' \in \langle [a] \rangle$. Since $c \in \text{supp}(\beta)$, we know $c \neq b^{\pm 1}$. We note that if d is $b^{\pm 1}$, then it must be that $w'u$ contains a generator not in $\text{st}(b)$. If w' is trivial, then u contains elements adjacently dominated by a , which cannot be in $\text{st}(b)$ since a is not adjacent to b . If w' is nontrivial, then w' contains elements of $[a]$ which are not adjacent to b . Further $w'u$ contains no elements of $\text{supp}(\beta)$ because β fixes $\text{st}(a)$. So whether $d = b^{\pm 1}$ or not, the initial segment of $cw'ud$ contains a syllable with respect to $[b]$ starting with c (in $\text{supp}(\beta)$) and ending with an element not in $\text{supp}(\beta)$. Such a syllable increases in length by one under β . Then summing over all syllables in T' that increase in length under α_2^{-1} , we have

$$L_2 \leq |v|L_1.$$

Summing these gives us

$$|v|S_1 + L_2 \leq |v|L_1 + S_2,$$

in other words

$$|v|(|W| - |\beta \cdot W|) = |v|(S_1 - L_1) \leq S_2 - L_2 = |\alpha_1 \cdot W| - |\alpha_3 \cdot W|.$$

We can rewrite this as

$$|W| - |\alpha_3 \cdot W| \geq |v|(|W| - |\beta \cdot W|) + |W| - |\alpha_1 \cdot W|.$$

Since (W, α_1, β) is a peak, we know $|W| - |\beta \cdot W| \geq 0$ and $|W| - |\alpha_1 \cdot W| \geq 0$ with at least one strict. In either case, it follows that $|\alpha_3 \cdot W| < |W|$. As explained above, this means (W, α_3, β) is a peak that Lemma 6.13 lowers, and by attaching α_2^{-1} to the lowering factorization, we get a peak-lowering factorization for the peak (W, α_1, β) .

This proves the first case. Now we move to a slightly more general case.

Case. *The proposition is true for general $\alpha \in \Omega_{[a]}$, $\beta \in \Omega_{[b]}$ with $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$. (We are still assuming a does not dominate b and b does not dominate a .)*

If $\alpha|_{[b]}$ and $\beta|_{[a]}$ are both the identity, then Lemma 6.13 lowers the peak. So we assume without loss of generality that $\alpha|_{[b]}$ is not the identity. Lemma 6.4 implies that β is a long-range automorphism in this case. If we also assume that $\beta|_{[a]}$ is not the identity, then α is also a long-range automorphism, and therefore Theorem 4.6 applies to lower the peak. So we assume that $\beta|_{[a]}$ is the identity.

Since β is a long-range automorphism, we peak-reduce β with respect to W to get a factorization $\beta = \beta_k \cdots \beta_1$ by classic long-range Whitehead automorphisms. It is possible that $|\beta_i \cdots \beta_1 \cdot W| < |W|$ for some i in $1, \dots, k$. If this is the case, we lower the peak (W, α, β_1) by the method above and concatenate $\beta_k \cdots \beta_2$ to this factorization to get a peak-lowering factorization for $\beta\alpha^{-1}$.

Otherwise, since $\beta_k \cdots \beta_1$ is peak reduced, we have $|\beta \cdot W| = |W|$, $|\alpha \cdot W| < |W|$ and $|\beta_1 \cdots \beta_i \cdot W| = |W|$ for $i = 1, \dots, k$. We prove the claim by induction on the length k of the factorization. We let α_1, α_2 and α_3 have the same meanings as above, with β_1 taking the role of β . Then $\beta_1 \alpha_1^{-1} = \alpha_3^{-1} \beta_1 \alpha_2^{-1}$ is a peak-lowering factorization for the peak (W, α_1, β_1) , using Lemma 6.13 as explained above. However, if $k > 1$, then we have a new peak $(\beta_1 \cdot W, \alpha_3, \beta_k \cdots \beta_2)$. The inductive hypothesis implies that this new peak can be lowered. Concatenating $\alpha_3^{-1} \beta_1$ onto the peak-lowering factorization for $(\beta_1 \cdot W, \alpha_3, \beta_k \cdots \beta_2)$ gives us a peak-lowering factorization for (W, α_1, β) , which gives us a peak-lowering factorization for the original peak using Lemma 6.14. This proves the second case.

General case. Finally we prove the proposition as stated. We have $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$. If α does not fix b , then since a does not dominate b , we know that α conjugates b by a nontrivial element of $\langle [a] \rangle$. Then by Lemma 6.14, we may replace α by its product with an inner automorphism and assume that α fixes b . Similarly, we assume that β fixes a .

Claim. *If $c^{\pm 1}$ is in $\text{supp}(\alpha) \cap \text{supp}(\beta)$, then either both a and b non-adjacently dominate c , or else the connected component Y of c in $\Gamma \setminus \text{st}(a)$ is also a connected component of $\Gamma \setminus \text{st}(b)$ and $Y \cup Y^{-1} \subset \text{supp}(\alpha) \cap \text{supp}(\beta)$.*

If $c^{\pm 1}$ is in $\text{supp}(\alpha) \cap \text{supp}(\beta)$, then one possibility is that a dominates c . If b were adjacent to c , it would imply that a is adjacent to b , which is not the case. So b is not adjacent to c . If b does not also dominate c , then there is a vertex adjacent to a and c but not adjacent to b . This means that a and c are in the same component of $\Gamma \setminus \text{st}(b)$; then $\beta(a) = a$ implies that $c, c^{-1} \notin \text{supp}(\beta)$ by Lemma 6.3. This is a contradiction, so we see that both a and b dominate c . Then if c were adjacent to a , it would imply that a and b are adjacent, so c is also not adjacent to a .

Next we suppose that c in $\text{supp}(\alpha) \cap \text{supp}(\beta)$ is in a connected component Y of $\Gamma \setminus \text{st}(a)$ with at least two vertices. First of all, b is not adjacent to c , since otherwise α would not fix b (by Lemma 6.3, since b would be in Y). Further, it cannot be the case that b dominates c —if it did, b would be adjacent to a vertex of Y other than c and therefore b would be in Y and α

would not fix b . So let Y' be the component of c in $\Gamma \setminus \text{st}(b)$. Suppose $d \in Y$. Then there is a path from d to c outside of $\text{st}(a)$. If this path intersects $\text{st}(b)$, then b would be in Y , which is impossible. So the path from d to c is outside $\text{st}(b)$, meaning that d is in Y' . Since d was arbitrary $Y \subset Y'$. By a parallel argument, $Y' \subset Y$, and therefore they are equal. This proves the claim.

We prove the proposition by induction on the cardinality of $\text{supp}(\alpha) \cap \text{supp}(\beta)$. The base case is that $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$, which we covered in a previous case.

Now we work the inductive step. Since $\text{supp}(\alpha) \cap \text{supp}(\beta)$ is nonempty, we select an element c from it. First we assume that both a and b non-adjacently dominate c . Fix a graphically reduced representative for W . Let n_a denote the number of instances of c or c^{-1} in the representative in a subword cud or duc^{-1} with $u \in \langle \text{st}(a) \rangle$ and $d \in [a]^{\pm 1}$ and let n_b be defined similarly with a replaced by b . Let n' count the number of remaining instances of c or c^{-1} in the representative (those that are counted by neither n_a or n_b). Let v_1 and v_2 be in $\langle [a] \rangle$ with $\alpha(c) = v_1 c v_2$ and let u_1 and u_2 be in $\langle [b] \rangle$ with $\beta(c) = u_1 c u_2$. Let α_c be the element of $\Omega_{[a]}$ with $\alpha_c(d) = \alpha(d)$ for all $d \neq c$, and such that if $\alpha(c) = v_1 c v_2$ then $\alpha_c(c) = v_1 c$. Let β_c be defined similarly, with $\beta(c) = u_1 c$. We note that α_c and β_c can be expressed as products of Laurence generators and are well-defined automorphisms.

The decomposition of W into syllables with respect to $[a]$ coming from our representative, with Lemma 6.6, tells us the following: at each instance of c counted by n_b or n' , α increases the length of W by $|v_2|$ more than α_c does, and at each instance of c counted by n_a , α may decrease the length of W by up to $|v_2|$ more than α does or increase it by up to $|v_2|$ more. Specifically,

$$|v_2|(n' + n_b - n_a) \leq |\alpha \cdot W| - |\alpha_c \cdot W| \leq |v_2|(n' + n_b + n_a).$$

Similarly for β :

$$|u_2|(n' + n_a - n_b) \leq |\beta \cdot W| - |\beta_c \cdot W| \leq |u_2|(n' + n_a + n_b).$$

Of course, since $c \in \text{supp}(\alpha) \cap \text{supp}(\beta)$, we know $|v_2| > 0$ and $|u_2| > 0$.

Suppose both $|\alpha_c \cdot W| \geq |\alpha \cdot W|$ and $|\beta_c \cdot W| \geq |\beta \cdot W|$. Then $n' + n_b - n_a \leq 0$ and $n' + n_a - n_b \leq 0$. Summing these, we see that $n' \leq 0$, but since n' is a nonnegative integer, $n' = 0$. Then $n_b = n_a$, and in this case we have both $|\alpha_c \cdot W| = |\alpha \cdot W|$ and $|\beta_c \cdot W| = |\beta \cdot W|$. By the definition of a peak, either $|\alpha \cdot W| < |W|$ or $|\beta \cdot W| < |W|$. So in this case we have either $|\alpha_c \cdot W| < |W|$ or $|\beta_c \cdot W| < |W|$.

Otherwise, we have either $|\alpha_c \cdot W| < |\alpha \cdot W|$ or $|\beta_c \cdot W| < |\beta \cdot W|$. Since the definition of a peak ensures that both $|\alpha \cdot W| \leq |W|$ and $|\beta \cdot W| \leq |W|$, in this case we also have either $|\alpha_c \cdot W| < |W|$ or $|\beta_c \cdot W| < |W|$.

Without loss of generality we suppose $|\alpha_c \cdot W| < |W|$. Then (W, α_c, β) is a peak that can be lowered by the inductive hypothesis. So we append the element $(\alpha_c \alpha^{-1}) \in \Omega_{[a]}$ onto a peak-lowering factorization for (W, α_c, β) to get a peak-lowering factorization for (W, α, β) .

The case where $c \in \text{supp}(\alpha) \cap \text{supp}(\beta)$ is in a connected component of $\Gamma \setminus \text{st}(a)$ is similar, with n_a , n_b and n' defined similarly. \square

Lemma 6.16. *Suppose a and b are vertices of Γ with $[a] \neq [b]$ and $\alpha \in \Omega_{[a]}$ and $\beta \in \Omega_{[b]}$ are classic long-range Whitehead automorphisms such that: $\alpha(b) = b$, $\text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset$ and $\beta(a)$ is a or ab . Then $\alpha\beta\alpha^{-1}$ is a classic long-range Whitehead automorphism in $\Omega_{[b]}$ and for any tuple W of cyclic words, we have*

$$|W| - |\beta \cdot W| = |\alpha \cdot W| - |\alpha\beta \cdot W|.$$

Proof. The fact that $\alpha\beta\alpha^{-1} \in \Omega_{[b]}$ is relation R4 from Section 2.4 in Day [3], with α being $(A, a)^{-1}$ and β being (B, b) in that statement. The equation of differences of lengths is essentially Sublemma 3.21 from Day [3], applied twice. Near the end of the proof of that statement, the equation above appears as an inequality (α and α^{-1} are switched and the terms rearranged, but it is the same assertion). The same inequality with W replaced by $\alpha \cdot W$ and α replaced by α^{-1} is the inequality in the reverse direction, proving the equation. \square

Proposition 6.17. *Suppose a is not adjacent to b , b dominates a and a does not dominate b . Then the peak (W, α, β) can be lowered.*

Proof. More precisely, we prove:

Claim. *Suppose a is not adjacent to b , b dominates a and a does not dominate b . Then the peak (W, α, β) has a peak-lowering factorization such that each automorphism in the factorization is in $\Omega_{[b]}$ or is a long-range automorphism in $\Omega_{[a]}$.*

Since a does not dominate b , we know that α acts on $[b]$ by conjugating it by a fixed element of $\langle [a] \rangle$. Then by Lemma 6.14, we can replace α by its composition with an inner automorphism from $\Omega_{[a]}$ and assume that α fixes $[b]$. We also note that since b non-adjacently dominates a , a does not adjacently dominate anything and therefore α is long-range and $[a] = \{a\}$ (see Lemma 6.4).

We induct on $|\alpha|_O$, the length of the outer class of α as a product of Laurence generators. We use as a base case the case that α is a classic long-range Whitehead automorphism, which will certainly be true if $|\alpha|_O$ is one. Before proving this case, we prove the inductive step. We peak-reduce α with respect to W by Theorem 4.6. We select the first automorphism α_1 out of such a peak-reducing factorization. Let $\alpha_2 = \alpha\alpha_1^{-1}$. Then α_1 will be a classic long-range Whitehead automorphism with multiplier $a^{\pm 1}$, $|\alpha_2|_O < |\alpha|_O$, and $|\alpha_1 \cdot W| \leq |W|$, with the inequality being strict if $|\alpha \cdot W| < |W|$. (Specifically, we factor α as a product of transvections and partial conjugations with multipliers $a^{\pm 1}$ and apply the algorithm from Day [3]; at each step, the algorithm merges, commutes or splits the automorphisms in the factorization, and the resulting peak-reduced factorization consists entirely of long-range automorphisms with multipliers in $\{a, a^{-1}\}$.) Then (W, α_1, β) is a peak. We apply the claim and get a peak-lowering factorization of $\beta\alpha_1^{-1}$. If the last term in the factorization is a long-range automorphism $\alpha' \in \Omega_{[a]}$, then $|\alpha'^{-1}\alpha_1 \cdot W| < |W|$ by the definition of a peak-lowering factorization. We peak-reduce $\alpha_2\alpha'$ with respect to $|\alpha'^{-1}\alpha_1 \cdot W|$ using Theorem 4.6. If we replace the element α' in our factorization of $\beta\alpha_1^{-1}$ with this peak-reduced factorization of $\alpha_2\alpha'$, then the result is a peak-lowering factorization of $\beta\alpha^{-1}$. Otherwise the last term in the factorization of $\beta\alpha_1^{-1}$ is an automorphism $\beta' \in \Omega_{[b]}$. Then $(\alpha_1 \cdot W, \alpha_2, \beta')$ is a peak satisfying the hypotheses of the claim, but with $|\alpha_2|_O < |\alpha|_O$. Then by induction, we apply the claim and get a peak-lowering factorization of $\beta'\alpha_2^{-1}$. We replace β' in the factorization of $\beta\alpha_1^{-1}$ with the entire factorization of $\beta'\alpha_2^{-1}$ to get a peak-lowering factorization of $\beta\alpha^{-1}$.

Now we prove the claim in the base case: with the additional hypothesis that α is a classic long-range Whitehead automorphism with multiplier a (the case where the multiplier is a^{-1} is similar). Let β'' be the short-range part of β ; in other words let β'' equal β on $\text{st}(b)$ and let β fix all other vertices. Then β'' will be a product of short-range transvections and inversions and will be a well defined automorphism. Let β' be the difference, the long-range part, so that $\beta = \beta'\beta''$. Everything in $\text{supp}(\beta'')$ is adjacent to b and dominated by b , or equal to b . Since a is not adjacent to b , a cannot be adjacent to anything in $\text{supp}(\beta'')$ (otherwise domination would force a to be adjacent to b). Since α fixes b , α fixes the entire connected component of b in $\Gamma \setminus \text{st}(a)$. Therefore $\text{supp}(\beta'') \cap \text{supp}(\alpha) = \emptyset$. Since a is not adjacent to b , by definition, β'' fixes $[a]$. Then by Lemma 4.7 and Proposition 6.7, we know α commutes with β'' and

$$|W| - |\alpha \cdot W| = |\beta'' \cdot W| - |\alpha\beta'' \cdot W|.$$

So $\beta\alpha^{-1} = \beta'\alpha^{-1}\beta''$.

Since we are about to start a step that we will repeat, we relabel α as α' and $\beta'^{-1}\beta \cdot W$ as W' . Let α'' denote the trivial automorphism at first. Then we have $\beta\alpha^{-1} = \beta'\alpha'^{-1}\beta''\alpha''^{-1}$. We process this factorization into a better one using an algorithm with a loop. If $|W'| < |W|$, then we do not enter the loop and instead skip ahead to the next step. Otherwise (W', α', β') is a peak:

$$|\beta' \cdot W'| = |\beta \cdot W| \leq |W| \leq |W'|$$

and

$$|W'| - |\alpha' \cdot W'| = |\beta'' \cdot W| - |\alpha'\beta'' \cdot W| = |W| - |\alpha' \cdot W| \geq 0,$$

with one of these inequalities strict because (W, α, β) is a peak.

Claim. *There is an algorithm to iteratively process the factorization $\beta\alpha^{-1} = \alpha''^{-1}\beta''\alpha'^{-1}\beta'$, such that at the beginning of each loop we have*

- (W', α', β') is a peak, where $W' = \beta'^{-1}\beta \cdot W$,
- α' is a non-inner long-range classic Whitehead automorphism,
- β' is a non-inner automorphism,
- $\alpha'^{-1}\beta''\alpha'$ is in $\Omega_{[b]}$ and

$$|W'| - |\alpha' \cdot W'| = |\beta''^{-1} \cdot W'| - |\beta''^{-1}\alpha' \cdot W'|.$$

With each repeat of the loop, we strictly decrease $|W'|$, or we do not increase $|W'|$ but strictly decrease one of $|\alpha'|_O$ or $|\beta'|_O$ while leaving the other fixed. At the end of an iteration, either the new $|W'|$ satisfies $|W'| < |W|$ and we exit the loop, or we repeat the loop (and satisfy the conditions for beginning the loop with the relabeled terms).

We have already shown that if we enter the loop at all, then the beginning conditions are satisfied for the first iteration. Now we explain the algorithm. By construction, β' is a long-range automorphism. Using Theorem 4.6, we peak-reduce β' with respect to W' and consider the first automorphism β_0 in the resulting peak-reduced factorization of β' . This β_0 is a classic long-range Whitehead automorphism with multiplier d^ϵ for some $d \in [b]$ and $\epsilon = \pm 1$. Since $|\beta' \cdot W'| \leq |W'|$ (because $|\beta' \cdot W'| = |\beta \cdot W| \leq |W|$), it follows from the definition of a peak-reduced factorization that $|\beta_0 \cdot W'| \leq |W'|$. Then (W', α', β_0) is a peak.

If $\text{supp}(\alpha') \cap \text{supp}(\beta_0) = \emptyset$ and $a \notin \text{supp}(\beta_0)$, then Lemma 6.16 applies. In particular, $\alpha' \beta_0 \alpha'^{-1} \in \Omega_{[b]}$ and

$$|W'| - |\alpha' \cdot W'| = |\beta_0 \cdot W'| - |\alpha' \beta_0 \cdot W'|.$$

We replace β' with $\beta_0^{-1} \beta'$ and replace β'' with $\beta'' \alpha' \beta_0 \alpha'^{-1}$. The new W' is $\beta_0 \cdot W'$. We go back to the beginning of the loop since the conditions to continue the loop are satisfied. If $\text{supp}(\alpha') \subset \text{supp}(\beta_0)$ and $a \in \text{supp}(\beta_0)$, we replace β_0 with its complement so that we have $\text{supp}(\alpha') \cap \text{supp}(\beta_0) = \emptyset$ and $a \notin \text{supp}(\beta_0)$. Then we apply the previous case. In these cases we have not increased $|W'|$ or the length of α' , but we have decreased the length of β' . If the new β' is inner, then the changes in length imply that the new W' is shorter than $\beta \cdot W$ and is therefore shorter than W .

Now we want $a \in \text{supp}(\beta_0)$; if this is not the case, we replace β_0 with its complement. This replacement does not change $\beta_0 \cdot W$, so (W', α', β_0) is still a peak. After performing such a replacement if necessary, we definitely have $a \in \text{supp}(\beta_0)$. Further, since we have just considered this case separately, we can assume that $\text{supp}(\alpha') \not\subset \text{supp}(\beta_0)$. Then we apply Corollary 6.9 to (W', α', β_0) and one of the following holds:

- there is a classic long-range Whitehead automorphism β_1 with multiplier $d^{-\epsilon}$ such that $|\beta_1 \cdot W'| < |W'|$ and $\text{supp}(\beta_1) \cap \text{supp}(\alpha) = \emptyset$ and β_1 fixes $[a]$, or
- there is a classic long-range Whitehead automorphism α_1 with multiplier a such that $|\alpha_1 \cdot W'| < |W'|$ and $\text{supp}(\alpha_1) \subset \text{supp}(\alpha') \cap \text{supp}(\beta_0)$.

If $\text{supp}(\alpha') \cap \text{supp}(\beta_0) = \emptyset$, then the first case holds.

If the second case holds, $\alpha' \alpha_1^{-1}$ commutes with β'' since $\text{supp}(\alpha' \alpha_1^{-1}) \subset \text{supp}(\alpha')$ and $\text{supp}(\alpha') \cap \text{supp}(\beta'') = \emptyset$, and since $\alpha' \alpha_1^{-1}$ fixes $[b]$ and β'' fixes $[a]$. Note that $\alpha_1 \neq \alpha$ since we assumed $\text{supp}(\alpha') \not\subset \text{supp}(\beta_0)$. So we relabel α_1 as α' and relabel α'' as $\alpha'' \alpha' \alpha_1^{-1}$. Then we still have the conditions we expect at the beginning of the loop. The new α' is nontrivial because $\alpha_1 \cdot W' \neq W'$ and (W', α', β') is still a peak because $|W'|$ is now strictly greater than $|\alpha' \cdot W'|$ and $|W'| - |\beta' \cdot W'|$ is unchanged. Then we repeat the loop. In this case we have not increased $|W'|$ or the length of β' but we have decreased the length of α' .

In the first case, Lemma 4.7 and Proposition 6.7 imply that α' and β_1 commute and that

$$|W'| - |\alpha' \cdot W'| = |\beta_1 \cdot W'| - |\alpha' \beta_1 \cdot W'|.$$

We relabel $\beta_1^{-1}\beta'$ as β' , and relabel $\beta_1\beta''$ as β'' . Note that we still have $\beta\alpha^{-1} = \beta'\alpha'^{-1}\beta''\alpha''^{-1}$. Further, we still have β'' fixing $[a]$ and $\text{supp}(\beta'') \cap \text{supp}(\alpha') = \emptyset$, since this was true of β_1 and the old β'' . We relabel W' as $\beta'^{-1}\beta \cdot W = \beta_1 \cdot W'$. We also note that the new $|W'|$ is strictly less than the old $|W'|$, since it is the image of the old W' under β_1 . If $|W'| < |W|$, then we exit the loop. Otherwise (W', α', β') is still a peak since α' shortens the new and old W' by the same amount. We repeat the loop with the newly relabeled automorphisms. If the new β' is inner, then the changes in length imply that the new W' is shorter than $\beta \cdot W$ and is therefore shorter than W .

So we leave the loop when $|W'| < |W|$. The sum of differences

$$|\alpha''^{-1}\alpha \cdot W| - |\alpha \cdot W| + |W'| - |\alpha' \cdot W'|$$

equals $|W| - |\alpha \cdot W|$. These two facts together are enough to deduce that $\beta\alpha^{-1} = \alpha''^{-1}\beta''\alpha'^{-1}\beta'$ is a peak-lowering factorization. \square

6.5 Peak reduction and stabilizer presentations

Finally, we use peak reduction to homotope a path in the complex Z from Section 4.4 in order to prove Proposition 4.9. We need the following.

Lemma 6.18. *Consider the set of conjugacy classes of A_Γ consisting of all conjugacy classes of length one represented by a positive generator and all conjugacy classes of length two represented by a product of two non-commuting generators. Let U be a tuple of conjugacy classes whose entries are this set of classes in some order. Then U is minimal length in its automorphism orbit.*

Further, if α is a generalized Whitehead automorphism sending one minimal-length representative of the orbit of U to another, then α is a permutation automorphism or an inner automorphism.

Proof. Without loss of generality, we suppose that U is ordered with the conjugacy classes of length one first and the classes of length two following. By peak reduction, to show that U is minimal length in its orbit, it is enough to show that no element of $\Omega_{[a]}$ can shorten it, for any a in Γ . Let T be a syllable decomposition of U ; we claim that the matrix $\nu(T)$ can be put in $G_{\mathbb{Q}}$ -normal form by permuting the rows. To see this we examine the matrix $\nu(T)$.

In the following, ϵ and δ are always in $\{1, -1\}$. For each b adjacent to or equal to a , the conjugacy class represented by b^ϵ maps to ϵr_b under

ν . For each b not adjacent to a , with a dominating b , the conjugacy class represented by b^ϵ is a single syllable that maps to $\epsilon r_b + \epsilon l_b$ under ν . For b not adjacent to a , with a not dominating b , the conjugacy class represented by b^ϵ is a single syllable that maps to 0 under ν (it maps to $r_Y - r_Y$ where Y is the component of b in $\Gamma \setminus \text{st}(a)$). For b and c both adjacent to or equal to a but not adjacent to each other, the conjugacy class of $b^\epsilon c^\delta$ maps to $\epsilon r_b + \delta r_c$. For b adjacent to or equal to a and c not adjacent to a or b , but with a dominating c , the conjugacy class of $b^\epsilon c^\delta$ maps to $\epsilon r_b + \delta r_c + \delta l_c$. For b adjacent to or equal to a and c in the component Y of $\Gamma \setminus \text{st}(a)$ (which has at least two vertices), the conjugacy class of $b^\epsilon c^\delta$ maps to ϵr_b . For b and c not adjacent to a , but with a dominating both b and c , the conjugacy class of $(bc)^\epsilon$ splits into two syllables, which map to $\epsilon(r_b + l_c)$ and $\epsilon(r_c + l_b)$, and the conjugacy class of $(bc^{-1})^\epsilon$ splits into two syllables, which map to $\epsilon(r_b - r_c)$ and $\epsilon(l_b - l_c)$. If b is not adjacent to a , but a dominates b , and c is in the connected component Y of $\Gamma \setminus \text{st}(a)$ (which has at least two elements), then the conjugacy class of $(bc^\delta)^\epsilon$ splits into two syllables, which map to $\epsilon(r_b - r_Y)$ and $\epsilon(l_b + r_Y)$. Finally, if b is in the component Y of $\Gamma \setminus \text{st}(a)$ and c is in the component Z of $\Gamma \setminus \text{st}(a)$ (with both having at least two elements), then the conjugacy class of $b^\epsilon c^\delta$ splits into two syllables that map to $r_Y - r_Z$ and $r_Z - r_Y$.

Let n denote the cardinality of $[a]$ and let k denote the number of basis elements of $Z_{[a]}$ other than the r_b for $b \in [a]$. First we verify that adding any linear combination of the last k rows of $\nu(T)$ to any of the first n does not simplify the matrix. If c does not commute with a and b is in $[a]$, then adding combinations of the row for r_c and l_c to the row for r_b may simplify the columns for $b^\epsilon c^\delta$, but any such action will make the column for c more complicated. By our hypothesis on U , the column for c precedes all the $b^\epsilon c^\delta$ columns of $\nu(T)$, and therefore none of these row moves can simplify $\nu(T)$. However, none of the columns for the other kinds of conjugacy classes in U (enumerated above) have a nonzero element in the first n rows and a nonzero element in the last k rows. This implies that adding linear combinations of the last k rows to the first n rows cannot simplify any column of $\nu(T)$ without making a previous column worse. Since the first $|X|$ columns of $\nu(T)$ include the n columns coming from the classes b for b in $[a]$, which map to r_b in each case, and since all other columns with a nonzero entry in the first n rows appear after the first $|X|$ columns, the matrix $\nu(T)$ can be put in $G_{\mathbb{Q}}$ -normal form by a permutation of the first n rows.

If an element of $\Omega_{[a]}$ shortens U , then all it can do is to delete some elements of $[a]$ from somewhere in U . Such a deletion would give us a tuple U' with a syllable decomposition T' differing only from T in some deletions

of elements of $[a]$. The corresponding matrix $\nu(T')$ would also essentially already be in $G_{\mathbb{Q}}$ -minimal form (up to a permutation), and therefore U and U' cannot be in the same orbit under $\Omega_{[a]}$ by Proposition 5.3.

To see the second part of the statement, we suppose that α is a generalized Whitehead automorphism sending U to a tuple of the same length. Of course, it is possible that α is a permutation automorphism. We suppose that it is not, so that α is in $\Omega_{[a]}$ for some a . By the form of $\nu(T)$, we know that the only way α can send $\nu(T)$ to itself is to add a linear combination of the bottom k rows of $\nu(T)$ to some rows in the top, where that linear combination has a trivial value. However, the only such linear combination that is trivial is the one corresponding to conjugation by some element: we add 0 times each r_b with b commuting with a , +1 times each r_c and -1 times each l_c with c not commuting with a and with a dominating c , and +1 times each r_Y row. This can be seen from the enumeration of columns above. \square

Proof of Proposition 4.9. Recall that Z is our alleged presentation complex for $\text{Aut}(A_{\Gamma})_W$ and we have an edge-loop p based at W that we wish to homotope so that its edge labels are only permutations and inner automorphisms.

We define a non-locally finite graph \widehat{Z} that is related to Z . The vertices of \widehat{Z} are the same as the vertices of Z . The edges are as follows: if W' is a vertex in \widehat{Z} and α is a generalized Whitehead automorphism in P or in $\Omega_{[a]}$ for some a with $|\alpha \cdot W'| = |W'|$, then there is an edge labeled by α from W' to $\alpha \cdot W'$. Of course Z^1 is a finite subgraph of \widehat{Z} .

Now we define a map ϕ from edge paths in \widehat{Z} to edge paths in Z^1 . For edges in $Z^1 \subset \widehat{Z}$, we define ϕ to be the identity. Suppose we have an edge in \widehat{Z} starting at a vertex W' that is not in Z^1 . Then this edge is labeled by an element α in $\Omega_{[a]}$ for some $a \in X$. Let S denote $(X \setminus [a])^{\pm 1} \setminus \text{supp}(\alpha)$. Then $\alpha \in \Omega_{[a], S}$. By the construction of Z , there is an edge starting at W' in Z^1 labeled by some $\gamma \in \Omega_{[a], S}$. Since $\alpha\gamma^{-1} \in (\Omega_{[a], S})_{W'}$, there is an edge path w in the loops at W' representing $\alpha\gamma^{-1}$ (this step involves a choice, but we make these choices once and for all and forget about them). Then ϕ sends the edge in \widehat{Z} starting at W' labeled with α to the concatenation of the path w with the edge labeled by γ . We extend ϕ to all edge paths in Z^1 by concatenation.

For any edge path in \widehat{Z} , the composition of edge labels gives us the same automorphism before and after applying ϕ . Further, the initial and terminal vertices of an edge path are the same before and after applying ϕ .

Our path p corresponds to a factorization $1 = \alpha_k \cdots \alpha_1$ such that the

intermediate images $\alpha_i \cdots \alpha_1 \cdot W$ have the same length as W and there is an edge labeled by α_{i+1} from $\alpha_i \cdots \alpha_1 \cdot W$ to $\alpha_{i+1} \cdots \alpha_1 \cdot W$ for each i . By Lemma 6.18, there is a tuple U of cyclic words, minimal length in its orbit, with the following property: any generalized Whitehead automorphism sending one minimal length image of U to another must be an inner automorphism or a permutation automorphism. We consider the intermediate images $\alpha_i \cdots \alpha_1 \cdot U$ of U and let m denote the maximum length of any of these. Let V be the concatenation of U with m copies of W . Our plan to prove the proposition is to peak-reduce the factorization $\alpha_k \cdots \alpha_1$ with respect to V . Peak reduction proceeds by finding peak-lowering substitutions to apply to peaks (subwords of length two) in the factorization. Since we have m copies of W in V , a peak-lowering substitution can never produce a new factorization that has an intermediate image of W longer than the original W —any increase in length in an image of W would have to be countered by a decrease in length in an image of U that is greater than the length of the longest intermediate image of U . Then if we peak-reduce $\alpha_k \cdots \alpha_1$ with respect to V , at each step the factorization remains the composition of labels on an edge path in Z . So to prove the proposition, it is enough to explain how the peak-lowering substitutions from the proof of Theorem 1.3 correspond to homotopies of the path p .

By inspecting Sections 6.3 and 6.4 above, we can see that the peak-lowering substitutions we use are from the following list:

- Factorization by classic Whitehead automorphisms: we take a long-range Whitehead automorphism from Ω and replace it with a product of classic long-range Whitehead automorphisms that follow an edge path in Z .
- Applying classic peak reduction to a peak between two classic long-range Whitehead automorphisms.
- Conjugating an automorphism across an inner automorphism as in Lemma 6.14.
- Applying a relation of length three between Whitehead automorphisms with the same multiplier, as in Lemma 6.11.
- Conjugating a permutation automorphism across another automorphism, as in Lemma 6.10.
- Applying a Steinberg relation from Lemma 4.7.

We proceed to show how applying any of the above rules to $\alpha_k \cdots \alpha_1$ corresponds to homotoping the path given by $\phi(\alpha_k \cdots \alpha_1)$ across some of the 2-cells in Z . Applying a substitution to $\alpha_k \cdots \alpha_1$ corresponds to homotoping $\phi(\alpha_k \cdots \alpha_1)$ across 2-cells in Z . In each of the following items, we indicate which cells we need.

- First we suppose we replace a long-range automorphism by a peak-reduced product of classic Whitehead automorphisms. We suppose we have a vertex W' in \widehat{Z} and a long-range generalized Whitehead automorphism γ labeling an edge in \widehat{Z} starting at W' . We know ϕ sends γ to a product $\gamma_2 \gamma_1$ labeling edges starting at W' in Z , where γ_i has the same multiplier and support as γ , for $i = 1, 2$, and where γ_1 fixes W' and $\gamma_2 \cdot W' = \gamma \cdot W'$. Both γ_1 and γ_2 are long-range automorphisms, so we can drag the path they follow across 2-cells in Z of type (C2) in order to replace $\gamma_2 \gamma_1$ with a product $\delta_l \cdots \delta_1$ of long-range classic Whitehead automorphisms tracing out an edge path in Z . We substitute this factorization back into $\alpha_k \cdots \alpha_1$, replacing γ , and homotope the image under ϕ across cells of type (C2).
- Now we suppose we must perform classic peak reduction. Suppose $\beta\gamma$ is a subword of length two of $\alpha_k \cdots \alpha_1$ and β and γ are classic long-range Whitehead automorphisms that we must peak-reduce with respect to V' , an intermediate image of V . Since V' is an intermediate image of V , it is the concatenation of U' with several copies of W' , where U' is an intermediate image of U and W' is an intermediate image of W . Since β and γ are classic, ϕ sends the edges with these labels in \widehat{Z} to edges with the same labels in Z , and W' is the vertex of Z that these edge are based at. Lowering this peak with respect to V' will produce a factorization that is peak-lowering with respect to U' , but which never increases the lengths of the intermediate images of W . Therefore lowering this peak with respect to V' will give us a factorization $\beta\gamma = \delta_l \cdots \delta_1$ that follows a path in Z . By Lemma 5.1 of Day [3], the substitution lowering this peak is an application of a relation between classic Whitehead automorphisms. Since these relations are all witnessed by 2-cells of type (C3), we can replace $\beta\gamma$ by $\delta_l \cdots \delta_1$ in $\alpha_k \cdots \alpha_1$ and homotope $\phi(\alpha_k \cdots \alpha_1)$ across cells of type (C3).
- Next we suppose that we have a vertex W' in \widehat{Z} and a generalized Whitehead automorphism $\beta \in \Omega_{[b]}$ for some $b \in X$, labeling an edge starting at W' and an inner classic Whitehead automorphism γ , and the next step in the peak reduction of $\alpha_k \cdots \alpha_1$ requires us to conjugate

β across γ . We know that β is $\beta_2\beta_1$, where $\beta_1 \in (\Omega_{[b]})_{W'}$ and $\beta_2 \in \Omega_{[b]}$ labels an edge starting at W' in Z^1 . We conjugate β_1 and β_2 across γ , which labels a loop in Z^1 . This amounts to a homotopy of $\phi(\alpha_k \cdots \alpha_1)$ across 2-cells of type (C4).

- Now we suppose $\alpha_k \cdots \alpha_1$ has a subword $\gamma\beta$ with $\beta, \gamma \in \Omega_{[b]}$ for some $b \in X$, and we need to replace $\gamma\beta$ by its product $\delta = \gamma\beta \in \Omega_{[b]}$ in the factorization as the next step in the peak reduction. Here β is an edge label on an edge originating at a vertex W' in \hat{Z} . Then $\phi(\beta)$ is $\beta_2\beta_1$, $\phi(\gamma)$ is $\gamma_2\gamma_1$, and $\phi(\delta)$ is $\delta_2\delta_1$, where all these new automorphisms are in $\Omega_{[b]}$, β_1 , γ_1 and δ_1 stabilize the vertices they originate at, and β_2 , γ_2 and δ_2 label edges in Z^1 . So $\phi(\alpha_k \cdots \alpha_1)$ will differ before and after the replacement as follows: before we have $\gamma_2\gamma_1\beta_2\beta_1$ and after we have $\delta_2\delta_1$. Pulling $\gamma_1\beta_2$ across cells of type (C5) will replace it with $\beta_2\gamma_1$, where the second γ_2 is a product of edge labels at W' instead of $\beta_2 \cdot W'$. Then we can homotope $\gamma_2\beta_2$ to $\delta_3\delta_1$ using a (C5) cell, where δ_3 is some product of edge labels at W' . Then we have homotoped $\gamma_2\gamma_1\beta_2\beta_1$ to a product of δ_2 with a product of edge labels at W' . At this point we only need to use some cells of type (C1) to finish this case.
- For this item we suppose $\alpha_k \cdots \alpha_1$ has a subword $\gamma\beta$ with $\gamma \in \Omega_{[a]}$ for some $a \in X$ and β a permutation automorphism. Suppose we need to replace this $\gamma\beta$ with $\beta\delta$, where $\delta = \beta^{-1}\gamma\beta \in \Omega_{[\beta^{-1}(a)]}$. Suppose γ in this sequence is the label on an edge originating at a vertex W' . We know $\phi(\gamma)$ is some $\gamma_2\gamma_1$ where $\gamma_1 \in (\Omega_{[a]})_{W'}$ and γ_2 is the label on an edge in Z originating at W' , and $\phi(\delta)$ is some $\delta_2\delta_1$ with $\delta_1 \in (\Omega_{[\beta^{-1}(a)]})_{\beta^{-1}\gamma \cdot W'}$ and δ_2 is the label on an edge in Z originating at $\beta^{-1}\gamma \cdot W'$. We know γ_1 is represented by some edge path in the loops at W' . For each edge in the sequence, we can conjugate β across the given edge label, thereby pushing our edge path across a 2-cell of type (C6). Then we conjugate β across γ_2 , getting $\beta\delta_2$ together with some product of loops at $\beta^{-1}\gamma \cdot W'$. We use disks of type (C1) to rewrite the remaining sequence as $\beta\delta_2\delta_1$. This homotopes the image of $\alpha_k \cdots \alpha_1$ under ϕ from before the replacement to the image after the replacement.
- Now we suppose we have $\beta\gamma$ in $\alpha_k \cdots \alpha_1$ with $\beta \in \Omega_{[b]}$ and $\gamma \in \Omega_{[a]}$ satisfying the hypotheses of Lemma 4.7. Suppose β is an edge label originating from the vertex W' and to peak-reduce $\alpha_k \cdots \alpha_1$ we need to conjugate γ across β . Let $\delta = \gamma^{-1}\beta\gamma$, which is in $\Omega_{[b]}$ by Lemma 4.7. Then ϕ sends the γ originating at $\gamma^{-1} \cdot W'$ to some $\gamma_2\gamma_1$ and sends the γ

originating at $\delta\gamma^{-1} \cdot W'$ to some $\gamma_4\gamma_3$, where these automorphisms are in $\Omega_{[a]}$, we know γ_2 and γ_4 label edges leaving $\gamma^{-1} \cdot W'$ and $\delta\gamma^{-1} \cdot W'$ in Z , and γ_1 and γ_3 fix $\gamma^{-1} \cdot W'$ and $\delta\gamma^{-1} \cdot W'$ respectively. Let $S = (X \setminus [a])^\pm \setminus \text{supp}(\gamma)$. Then $\gamma \in \Omega_{[a],S}$, and by the construction of ϕ , we know γ_i is as well, for $i = 1, 2, 3, 4$. Similarly, ϕ sends β to $\beta_2\beta_1$ where β_2 labels an edge of Z starting at W' , β_1 fixes W' , $\beta_i \in \Omega_{[b]}$ and $\text{supp}(\beta_i) \subset \text{supp}(\beta)$ for $i = 1, 2$, and ϕ sends δ to $\delta_2\delta_1$ where δ_2 labels an edge of Z starting at $\gamma^{-1} \cdot W'$, δ_1 fixes $\gamma^{-1} \cdot W'$, $\delta_i \in \Omega_{[b]}$ and $\text{supp}(\delta_i) \subset \text{supp}(\delta)$ for $i = 1, 2$. These conditions mean that the edges labeled with β_2 and γ_2^{-1} leaving W' in Z satisfy the hypotheses for Lemma 4.7, and therefore there is a 2-cell of type (C7) for this relation. We can move β_1 out of the way using cells of type (C5). Then we slide the edge path across our cell of type (C7), and replace $\gamma_2\beta_2$ with $\gamma_2w_2w_1\delta_2$, where w_2 is an edge loop label sequence representing an element of $(\Omega_{[a]})_{\delta\gamma^{-1} \cdot W'}$ and w_1 is an edge loop label sequence representing an element of $(\Omega_{[b]})_{\delta\gamma^{-1} \cdot W'}$. Again we use 2-cells of type (C5) to move w_1 to precede δ_2 . We note that the edge labels appearing in w_1 and in γ_1 satisfy the hypotheses of Lemma 4.7, and so by a sequence of slides across 2-cells of type (C7) we can move one past the other. Similarly, the labels in γ_1 together with δ_2 satisfy the hypothesis of Lemma 4.7, and we can slide γ_1 across δ_2 . Then δ_1 and w_1 both represent the same element of $(\Omega_{[b]})_{\delta\gamma^{-1} \cdot W'}$ and we can use 2-cells of type (C1) homotope one to the other. Similarly we homotope w_2 to $\gamma_3\gamma_1$. So we have homotoped $\beta_2\beta_1\gamma_2\gamma_1$ to $\gamma_4\gamma_3\delta_2\delta_1$ relative to endpoints, so we have homotoped the image of $\alpha_k \cdots \alpha_1$ under ϕ before this move to the image after this move.

So we peak-reduce $\alpha_k \cdots \alpha_1$ with respect to V and homotope $\phi(\alpha_k \cdots \alpha_1)$ at each step, so our edge path describes an edge loop in \widehat{Z} that sends each intermediate image of U to one of the same length as U . Then each α_i is an inner automorphism or a permutation automorphism by Lemma 6.18. \square

Acknowledgements

This research was supported by an NSF grant, award number DMS-1206981. I am grateful to all the people with whom I had conversations about this research, including Matt Clay, Michael Day, Michael Geline, Mark Johnson and Yo'av Rieck. My mother Rose Day has made important contributions to all my work, indirectly, in no small amount.

References

- [1] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125:141–158, 2007.
- [2] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.
- [3] Matthew B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geom. Topol.*, 13(2):817–855, 2009.
- [4] Matthew B. Day. Symplectic structures on right-angled Artin groups: between the mapping class group and the symplectic group. *Geom. Topol.*, 13(2):857–899, 2009.
- [5] Susan Hermiller and John Meier. Algorithms and geometry for graph products of groups. *J. Algebra*, 171(1):230–257, 1995.
- [6] Michael R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc. (2)*, 52(2):318–334, 1995.
- [7] Hai-Ning Liu, C. Wrathall, and Kenneth Zeger. Efficient solution of some problems in free partially commutative monoids. *Inform. and Comput.*, 89(2):180–198, 1990.
- [8] J. McCool. A presentation for the automorphism group of a free group of finite rank. *J. London Math. Soc. (2)*, 8:259–266, 1974.
- [9] J. McCool. On Nielsen’s presentation of the automorphism group of a free group. *J. London Math. Soc. (2)*, 10:265–270, 1975.
- [10] James McCool. Some finitely presented subgroups of the automorphism group of a free group. *J. Algebra*, 35:205–213, 1975.
- [11] John Milnor. *Introduction to algebraic K-theory*. Princeton University Press, Princeton, N.J., 1971. Annals of Mathematics Studies, No. 72.
- [12] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin, 1980. Translated from the French by John Stillwell.
- [13] Herman Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.

- [14] Leonard VanWyk. Graph groups are biautomatic. *J. Pure Appl. Algebra*, 94(3):341–352, 1994.
- [15] J. H. C. Whitehead. On equivalent sets of elements in a free group. *Ann. of Math. (2)*, 37(4):782–800, 1936.
- [16] C. Wrathall. Free partially commutative groups. In *Combinatorics, computing and complexity (Tianjing and Beijing, 1988)*, volume 1 of *Math. Appl. (Chinese Ser.)*, pages 195–216. Kluwer Acad. Publ., Dordrecht, 1989.